

GEOMETRIC FILTRATIONS OF CLASSICAL LINK CONCORDANCE

JAMES CONANT, ROB SCHNEIDERMAN, AND PETER TEICHNER

ABSTRACT. This paper studies *grope* and *Whitney tower* filtrations on the set of concordance classes of classical links in terms of *class* and *order* respectively. Using the tree-valued intersection theory of Whitney towers, the associated graded quotients are shown to be finitely generated abelian groups under a well-defined connected sum operation.

Twisted Whitney towers are also studied, along with a corresponding quadratic enhancement of the intersection theory for framed Whitney towers that measures Whitney-disk framing obstructions. The obstruction theory in the framed setting is strengthened, and the relationships between the twisted and framed filtrations are described in terms of exact sequences which show how higher-order Sato-Levine and higher-order Arf invariants are obstructions to framing a twisted Whitney tower.

The results from this paper combine with those in [9, 10, 11] to give a classifications of the filtrations; see our survey [8] as well as the end of the introduction below. UPDATE: The results of this paper have been completely subsumed into the paper *Whitney tower concordance of classical links* [13].

1. INTRODUCTION

Several key theorems and conjectures in low-dimensional topology can be stated in terms of certain 2-complexes known as *gropes*. These are geometric embodiments of commutators of group elements, see e.g. [1, 4, 5, 6, 14, 15, 27, 28]. Gropes are built by gluing together surfaces along collections of embedded essential circles, and come in several different types with varying measures of complexity.

The *height* of a grope corresponds to the derived series and was used (in the presence of caps) in [14, 15] to formulate the main open problem for topological 4-manifolds: the 4-dimensional surgery and s-cobordism theorems for arbitrary fundamental groups are equivalent to a certain statement about capped gropes of height ≥ 2 . In the uncapped setting,

Key words and phrases. Whitney towers, gropes, link concordance, trees, higher-order Arf invariants, higher-order Sato-Levine invariants, twisted Whitney disks.

gropes of increasing height were used in [4] to define a new filtration of the knot concordance group.

In this paper, we shall study a more basic measure of the complexity, namely the *class* of a grope, which corresponds to commutator length and the lower central series (Figure 1). In the 3-dimensional setting it was used in [5, 6] to give a geometric interpretation of Vassiliev's finite type invariants of knots and (the first nonvanishing term of) the Kontsevich integral.




FIGURE 1. A class 4 grope.

In the 4-dimensional setting it is a fundamental open problem to determine under what conditions the components of a link in the 3-sphere bound class n gropes disjointly embedded in the 4-ball, and this paper will describe a program for computing the
Grope concordance filtration (by class)

$$(\mathbb{G}) \quad \cdots \subseteq \mathbb{G}_3 \subseteq \mathbb{G}_2 \subseteq \mathbb{G}_1 \subseteq \mathbb{G}_0 \subseteq \mathbb{L}$$

on the set $\mathbb{L} = \mathbb{L}(m)$ of framed links in S^3 with m components. Here $\mathbb{G}_n = \mathbb{G}_n(m)$ is defined to be the set of framed links that bound class $(n+1)$ framed gropes disjointly embedded in B^4 . The index shift provides compatibility with other filtrations introduced below and for the same reason, we *define* \mathbb{G}_0 to be the set of evenly framed links.

The intersection of all \mathbb{G}_n contains all slice links because a 2-disk is a grope of arbitrary large class. In fact, this filtration factors through link concordance; and we shall use this fact implicitly at various places.

Recall that the connected sum of links is not a well-defined operation because of the choices of connecting bands that are involved. As a consequence, there is no direct definition of the associated graded quotients of the filtration \mathbb{G}_n . However, one can define an equivalence relation on links by using the notion of *class n grope concordance* between two links. This is obtained by using framed gropes built on annuli connecting link components in $S^3 \times I$. We then define the *associated graded* $\mathbb{G}_n = \mathbb{G}_n(m)$ as the quotient of \mathbb{G}_n modulo grope concordance of class $n+2$. Our first result is the following corollary of Theorem 3 below.

Corollary 1. *For all $m, n \in \mathbb{N}$, the sets $\mathbb{G}_n(m)$ are finitely generated abelian groups, under a well-defined connected sum $\#$. Moreover, \mathbb{G}_n is the set of framed links $L \in \mathbb{G}_n$ modulo the relation that $[L_1] = [L_2] \in \mathbb{G}_n$ if and only if $L_1\# - L_2$ lies in \mathbb{G}_{n+1} , for some choice of connected sum $\#$. Here $-L$ is the mirror image of L with reversed framing.*

For example, $\mathbb{G}_0(m) \cong \mathbb{Z}^k$ where $k = m(m + 1)/2$ is the number of possible linking numbers and framings (on the diagonal) of a link with m components. This follows from the fact that disjointly embedded class 2 gropes in B^4 are framed surfaces which induce the zero framings on their boundary and show the vanishing of all linking numbers. We also note that $\mathbb{G}_1(1) \cong \mathbb{Z}_2$ is given by the Arf invariant, and that $\mathbb{G}_2(1) = \mathbb{G}_n(1)$ for all $n \geq 2$, by [22].

The proof of Corollary 1 can be most succinctly formulated by defining a sequence of finitely generated abelian groups $\mathcal{T}_n = \mathcal{T}_n(m)$ in terms of certain trees. These groups have previously appeared in the study of graph cohomology, Feynman diagrams and finite type invariants of links and 3–manifolds.

Definition 2. [25] In this paper, a *tree* will always refer to an oriented unitrivalent tree, where the *orientation* of a tree is given by cyclic orientations at all trivalent vertices. The *order* of a tree is the number of trivalent vertices. Univalent vertices will usually be labeled from the set $\{1, 2, 3, \dots, m\}$ corresponding to the link components, and we consider trees up to isomorphisms preserving these labelings. We define $\mathcal{T} = \mathcal{T}(m)$ to be the free abelian group on such trees, modulo the antisymmetry (AS) and Jacobi (IHX) relations shown in Figure 2. Since

$$\text{AS: } \begin{array}{c} \text{---} \\ | \\ \text{---} \end{array} + \begin{array}{c} \text{---} \\ | \\ \text{---} \end{array} = 0$$

$$\text{IHX: } \begin{array}{c} \text{---} \\ | \\ \text{---} \end{array} - \begin{array}{c} \text{---} \\ | \\ \text{---} \end{array} + \begin{array}{c} \text{---} \\ | \\ \text{---} \end{array} = 0$$

FIGURE 2. Local pictures of the *antisymmetry* (AS) and *Jacobi* (IHX) relations in \mathcal{T} . Here all trivalent orientations are induced from a fixed orientation of the plane, and univalent vertices possibly extend to subtrees which are fixed in each equation.

the AS and IHX relations are homogeneous with respect to order, \mathcal{T}

inherits a grading $\mathcal{T} = \oplus_n \mathcal{T}_n$, where $\mathcal{T}_n = \mathcal{T}_n(m)$ is the free abelian group on order n trees, modulo AS and IHX relations.

Then a construction similar to Cochran's iterated Bing-doubling construction for realizing Milnor invariants [2] (and equivalent to ‘simple clasper surgery along trees’ in the sense of Goussarov and Habiro) leads to our *realization map*

$$R_n : \mathcal{T}_n \rightarrow \mathsf{G}_n$$

and Theorem 3 says that this map is well-defined and onto. For example, if $i \neq j$ then R_0 sends the order zero tree $i — j$ to a disjoint union of an $(m - 2)$ -component unlink and a Hopf link with components numbered i and j , all components being zero-framed. If $i = j$, we get an m -component unlink where the i th component has framing 2 and all other components are zero-framed.

In fact, it will turn out that G_n is *isomorphic* to \mathcal{T}_n when n is even (and to a quotient by certain 2-torsion in the odd case) via a map that is essentially defined by forgetting the roots of the rooted trees which correspond to iterated commutators determined by the grope branches.

The graded groups G_n are reminiscent of the groups of knots arising in the theory of finite type invariants, first discovered by Goussarov [16], defined as a quotient of the monoid of knots by n -equivalence. These groups were later constructed more conceptually using clasps by both Habiro and Goussarov [18, 17]. For the case of string links with m -strands, in [7] we defined $G_n(m)$ to be the monoid of m -string links modulo (3-dimensional) grope cobordism of class $n + 1$, and the associated graded quotients are finitely generated abelian groups. As in the current paper, we defined a surjective realization map

$$\Phi : \mathcal{B}_n^g(m) \rightarrow G_n(m)/G_{n+1}(m)$$

where $\mathcal{B}_n^g(m)$ is the group of Feynman diagrams with grope degree n , and showed that rationally it is an isomorphism. Clearly there is a surjection $\rho : G_n(m)/G_{n+1}(m) \twoheadrightarrow \mathsf{G}_n(m)$, and indeed our realization maps R_n are the compositions $\rho \circ \Phi$, since graphs with loops give rise to null-concordant links and are therefore in the kernel of $\rho \circ \Phi$. (See also Remark 28 below.)

Although we could take the composition $\rho \circ \Phi$ as our definition of R_n , and prove the required properties by modifying the 3-dimensional methods developed for knots by the first and third authors in [5, 6], we take a different approach here which is directly 4-dimensional and significantly simpler.

1.1. The Whitney tower filtration. We consider a second filtration on the set of framed links with m components, defined by using (framed) *Whitney towers* in place of (framed) gropes. Just like gropes are iterated surface stages glued to each other in a specified way, a Whitney tower is constructed by adding layers of immersed framed Whitney disks that pair up the intersection points of lower stages (Figure 3 and Section 2).

Whitney towers also come with a measure of complexity, the *order*, which determines how often intersection points are paired up. For example, an order zero Whitney tower is just a union of framed immersed disks in B^4 bounded by a link in S^3 with the induced framing, and an order 1 Whitney tower would also have all intersections among the immersed disks paired by Whitney disks. We define $\mathbb{W}_n = \mathbb{W}_n(m)$ to be the set of framed m -component links that bound a Whitney tower of order n in B^4 . On the set of framed links, we thus get the *Whitney tower filtration* (by order)

$$(\mathbb{W}) \quad \cdots \subseteq \mathbb{W}_3 \subseteq \mathbb{W}_2 \subseteq \mathbb{W}_1 \subseteq \mathbb{W}_0 \subseteq \mathbb{L} = \mathbb{L}(m)$$




FIGURE 3. Part of a Whitney tower in 4-space (left), and shown with part of an associated tree (right). Whitney disks pair up transverse intersections of opposite sign.

As for gropes above, the equivalence relation of order $n+1$ *Whitney tower concordance* (see Section 3) then leads to the ‘associated graded’ $\mathbb{W}_n = \mathbb{W}_n(m)$.

Theorem 3. *For each m, n , there are surjective realization maps*

$$R_n = R_n(m) : \mathcal{T}_n(m) \twoheadrightarrow \mathbb{W}_n(m)$$

and the sets $\mathbb{W}_n(m)$ are finitely generated abelian groups under the well-defined operation of connected sum $\#$. Moreover, \mathbb{W}_n is the set of framed links $L \in \mathbb{W}_n$ modulo the relation that $[L_1] = [L_2] \in \mathbb{W}_n$ if and

only if $L_1\# - L_2$ lies in \mathbb{W}_{n+1} , for some choice of connected sum $\#$. Here $-L$ is the mirror image of L with reversed framing.

Note that \mathbb{W}_0 consists of links that are evenly framed because a component has even framing if and only if it bounds a framed immersed disk in B^4 . It is well known that the signed number of self-intersection points of such a framed disk equals twice the framing on the boundary. As a consequence, if a knot bounds the next type of Whitney tower (of order one) then all the intersections of the first stage (order zero) disk are paired up by Whitney disks and so the framing is zero. In fact, $\mathbb{W}_0(1) \cong 2\mathbb{Z}$, given by the framing; and $\mathbb{W}_1(1) \cong \mathbb{Z}_2$, detected by the Arf invariant.

The resemblance between the two realization maps above is no coincidence: It is a result of the second author in [21] that $\mathbb{G}_n = \mathbb{W}_n$ and $\mathbb{G}_n = \mathbb{W}_n$ by completely geometric constructions that lead from gropes to Whitney towers and back (Figure 4). This means that we




FIGURE 4. Controlled conversions between order n Whitney towers and class $(n+1)$ gropes are described in detail in [21].

can continue to study either filtration and in this paper, we decide to take the Whitney tower point of view. This choice is motivated by the fact that Whitney towers come with an obstruction theory which is accompanied by higher-order intersection invariants taking values in \mathcal{T} . As described in Section 2 and hinted at in the right side of Figure 3, any Whitney tower \mathcal{W} of order n has an order n intersection invariant $\tau_n(\mathcal{W}) \in \mathcal{T}_n$ which is defined by summing the trees pictured in Figure 3. Order n intersection invariants vanish on Whitney towers of order $n+1$, and the key step in the proof of Theorem 3 (given in Section 3) is the following “raising the order” result:

Theorem 4 ([7, 21, 25]). *A link bounds a Whitney tower \mathcal{W} of order n with $\tau_n(\mathcal{W}) = 0$, if and only if it bounds a Whitney tower of order $n+1$.*

This theorem follows from realizing all relations in \mathcal{T}_n geometrically by controlled maneuvers on Whitney towers. It was a very pleasant surprise that, say, the Jacobi identity (or IHX relation) has an incarnation as a certain sequence of moves on Whitney disks and their boundary-arcs in a Whitney tower [7]. It has the following important corollary whose proof is in Section 36.

Corollary 5. *Links L_0 and L_1 represent the same element of \mathbb{W}_n if and only if there exist order n Whitney towers \mathcal{W}_i in B^4 with $\partial\mathcal{W}_i = L_i$ and $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1) \in \mathcal{T}_n$.*

Our realization map R_n (and its surjectivity) can then be explained extremely well by the following commutative diagram:

$$(1) \quad \begin{array}{ccc} \mathfrak{W}_n & \xrightarrow{\delta} & \mathbb{W}_n \\ \tau_n \downarrow & & \downarrow \\ \mathcal{T}_n & \xrightarrow{R_n} & \mathbb{W}_n \end{array}$$

An element in the set \mathfrak{W}_n is an order n Whitney tower \mathcal{W} in B^4 whose boundary $\partial\mathcal{W}$ lies in \mathbb{W}_n by definition of our filtration. The map R_n then arises directly from above the diagram via Corollary 5 and the surjectivity of the intersection invariant τ_n . The proof of Theorem 3 will be completed via a Bing doubling (and internal band sum) construction applied to the Hopf link, showing that τ_n is surjective.

As a consequence of Theorem 3 together with results in [9, 10, 11], it turns out that $\tau_{2n}(\mathcal{W})$ only depends on the concordance class of the link $L = \partial\mathcal{W}$, and not the Whitney tower \mathcal{W} it bounds; and that the intersection invariants τ_{2n} induce inverses to the realization maps R_{2n} . In particular, we have following classification theorem for even orders:

Theorem 6 ([11]). *The maps $R_{2n} : \mathcal{T}_{2n} \rightarrow \mathbb{W}_{2n}$ are isomorphisms.*

To describe the situation in odd orders, we first introduce the *reduced* version $\tilde{\mathcal{T}}_{2n-1}$ of \mathcal{T}_{2n-1} by dividing out the *framing relations*. These relations are the images of homomorphisms

$$\Delta_{2n-1} : \mathbb{Z}_2 \otimes \mathcal{T}_{n-1} \rightarrow \mathcal{T}_{2n-1}$$

defined by sending an order $n - 1$ tree t to the sum of trees gotten by doubling the subtree adjacent to each univalent vertex of t (see Figure 18 and the precise definition in Section 5).

Theorem 7. *The odd order realization maps R_{2n-1} vanish on the image of Δ_{2n-1} and give surjections*

$$\tilde{R}_{2n-1} : \tilde{\mathcal{T}}_{2n-1} \twoheadrightarrow \mathbb{W}_{2n-1}$$

Moreover, if a link L bounds a Whitney tower \mathcal{W} of order $2n - 1$ with $\tau_{2n-1}(\mathcal{W}) \in \text{Im}(\Delta_{2n-1})$ then L bounds a Whitney tower of order $2n$.

Here the results of [9, 10, 11] also apply to the cases where n is *even*, implying that $\tilde{\tau}_{4k-1}(\mathcal{W}) \in \tilde{\mathcal{T}}_{4k-1}$ only depends on the link L and not the Whitney tower \mathcal{W} , and that $\tilde{\tau}_{4k-1}$ induces an inverse to the realization map \tilde{R}_{4k-1} . We get the following classification in these orders:

Theorem 8 ([11]). $\tilde{R}_{4k-1} : \mathcal{T}_{4k-1} \rightarrow \tilde{\mathcal{W}}_{4k-1}$ are isomorphisms.

The framing relations in odd orders arise from a subtle interplay between the IHX relations and certain framing obstructions associated to Whitney disks, as will be described below. As suggested by theorems 7 and 8, we believe that these relations do indeed capture all the relevant geometric aspects regarding the filtration. Setting $\tilde{\mathcal{T}}_{2n} := \mathcal{T}_{2n}$, we conjecture:

Conjecture 9. *The realization maps give $\mathcal{W}_n \cong \tilde{\mathcal{T}}_n$ for all n .*

As discussed below and described in detail in [9, 11], the affirmation of this conjecture is equivalent to the non-triviality of certain *higher-order Arf invariants* which complete the classification of \mathcal{W}_{4k-3} .

The big picture is best described by introducing the following generalization of framed Whitney towers.

1.2. The Twisted Whitney tower filtration. The proof of Theorem 7 given in Section 5 uses in addition to IHX type maneuvers another well known geometric move on a Whitney disk, namely the *boundary twist*. A *framed* Whitney disk changes into a *twisted* Whitney disk by this move, at the cost of creating a new interior intersection in the Whitney disk (Figure 13).

This motivated us to introduce yet another filtration of the set of links by looking at those links $\mathcal{W}_n^\omega = \mathcal{W}_n^\omega(m)$ with m components that bound *twisted* Whitney towers of order n (rather than *framed* Whitney towers as in the definition of \mathcal{W}_n). These are Whitney towers, except that certain Whitney disks are allowed to be twisted. We arrive at the *Twisted Whitney tower filtration* (by order)

$$(\mathcal{W}^\omega) \quad \cdots \subseteq \mathcal{W}_3^\omega \subseteq \mathcal{W}_2^\omega \subseteq \mathcal{W}_1^\omega \subseteq \mathcal{W}_0^\omega = \mathbb{L}$$

We refer to Section 3 for a precise definition, also of the associated graded $\mathcal{W}_n^\omega = \mathcal{W}_n^\omega(m)$; and to Section 2 for details on twisted Whitney towers, including the associated *twisted intersection invariant* $\tau_n^\omega(\mathcal{W}) \in \mathcal{T}_n^\omega = \mathcal{T}_n^\omega(m)$. The groups $\mathcal{T}_{2n-1}^\omega$ are defined as quotients of \mathcal{T}_{2n-1} by the subgroups generated by trees of the form

$$i \prec_J^J$$

where J is a subtree of order $n - 1$. These *boundary-twist relations* correspond to the intersections created by performing a boundary-twist on an order n Whitney disk. The groups \mathcal{T}_{2n}^∞ include certain “twisted” ∞ -trees representing framing obstructions on order n Whitney disks, which are not required to be framed in an order $2n$ twisted Whitney tower. In [11] we show that \mathcal{T} is the universal home for invariant symmetric bilinear forms on quasi-Lie algebras, and that \mathcal{T}_{2n}^∞ is the universal (symmetric quadratic) refinement of this form in order n .

Here and in the following the symbol ∞ represents a *twist*, and in particular does *not* stand for “infinity”.

Theorem 10. *For each m, n , there are surjective realization maps*

$$R_n^\infty = R_n^\infty(m) : \mathcal{T}_n^\infty(m) \twoheadrightarrow \mathbb{W}_n^\infty(m)$$

and the sets $\mathbb{W}_n^\infty(m)$ are finitely generated abelian groups under the well-defined connected sum $\#$ operation. Moreover, \mathbb{W}_n^∞ is the set of framed links $L \in \mathbb{W}_n^\infty$ modulo the relation that $[L_1] = [L_2] \in \mathbb{W}_n^\infty$ if and only if $L_1\# - L_2$ lies in \mathbb{W}_{n+1}^∞ , for some choice of connected sum $\#$. Here $-L$ is the mirror image of L with reversed framing.

As before, the key step in proving this result is the following criterion for *raising the order of a twisted Whitney tower*:

Theorem 11. *A link bounds a twisted Whitney tower \mathcal{W} of order n with $\tau_n^\infty(\mathcal{W}) = 0$ if and only if it bounds a twisted Whitney tower of order $n + 1$.*

Theorem 11 follows from the more general Theorem 30 in Section 2, which is proved in Section 4 using “twisted Whitney moves” (Lemma 42) as well as boundary-twists and a construction for “geometrically cancelling” twisted Whitney disks.

In the twisted setting we also have classifications of the associated graded groups of links in three out of four cases:

Theorem 12 ([9, 11]). *For $n \equiv 0, 1, 3 \pmod{4}$, the realization maps $R_n^\infty : \mathcal{T}_n^\infty \rightarrow \mathbb{W}_n^\infty$ are isomorphisms.*

In fact, in these orders the \mathbb{W}_n^∞ are isomorphic to image D_n of the first nonvanishing order n (length $n + 2$) Milnor link invariants [9, 11]. Theorem 10 (as well as theorems 6 and 8 above) depends in an essential way on our resolution [10] of a combinatorial conjecture of J. Levine involving a map from unrooted trees representing Whitney tower intersections to rooted trees representing iterated commutators determined by link longitudes. As described in [9, 10], the kernel of this map is generated by certain symmetric ∞ -trees that are not detected

by Milnor invariants. These trees are related to what we call *higher-order Arf invariants* of the twisted filtration in [9], and the conjectured non-triviality of these invariants is shown in [11] to be equivalent to the above Conjecture 9 as well as the following:

Conjecture 13. *The maps $R_{4k-2}^\infty : \mathcal{T}_{4k-2}^\infty \rightarrow W_{4k-2}^\infty$ are isomorphisms.*

1.3. Framed versus twisted Whitney towers. There is a surprisingly simple relation between twisted and framed Whitney towers of various orders that's very well expressed in terms of the following result:

Proposition 14. *For any $n \in \mathbb{N}$, there is an exact sequence*

$$0 \longrightarrow W_{2n} \longrightarrow W_{2n}^\infty \longrightarrow W_{2n-1} \longrightarrow W_{2n-1}^\infty \longrightarrow 0$$

where all maps are induced by the identity on the set of links.

To see why this is true, observe that it is easier to find a twisted Whitney tower than a framed one and hence there is a natural inclusion $W_n \subseteq W_n^\infty$ between our two filtrations (W) and (W^∞) on the set of links. Moreover, by Definition 20 we actually have

$$W_{2n-1}^\infty = W_{2n-1}$$

because we found that a twisted Whitney disk of order n is most naturally associated to an order $2n$ Whitney tower, compare Remark 22. One then needs to show that indeed $W_{2n}^\infty \subseteq W_{2n-1}$, which is accomplished in Lemma 43 of Section 4 using boundary-twists and twisted Whitney moves. Finally, by the ‘Moreover’ parts of Theorems 3 and 10, we may say that

$$W_n = W_n / W_{n+1} \quad \text{and} \quad W_n^\infty = W_n^\infty / W_{n+1}^\infty$$

which implies the exact sequence in Proposition 14.

If our above conjectures above hold, then for every n the various realization maps should lead to the analogous exact sequence for our groups defined by trees. The following result shows that this is indeed the case and even gives more information on kernels (respectively cokernels). It is best expressed in terms of the *free quasi-Lie algebra* on m generators:

$$L' = L'(m) = \bigoplus_{n \in \mathbb{N}} L'_n$$

Definition 15 (compare [20]). The abelian group $L'_{n+1} = L'_{n+1}(m)$ is generated by order n (trivalent, oriented) *rooted* trees, each having a specified univalent vertex, labeled as *root*, and all other univalent vertices labelled by elements of $\{1, \dots, m\}$, modulo the AS and IHX relations of Figure 2.

Here the prefix ‘quasi’ reflects the fact that, although the IHX relation corresponds to the Jacobi identity via the usual identification of rooted trees with non-associative brackets, the usual Lie algebra self-annihilation relation $[X, X] = 0$ does *not* hold in \mathcal{L}' . It is replaced by the weaker anti-symmetry (AS) relation $[Y, X] = -[X, Y]$.

Theorem 16 ([11]). *For any $m, n \in \mathbb{N}$, there are short exact sequences*

$$0 \longrightarrow \mathcal{T}_{2n} \longrightarrow \mathcal{T}_{2n}^\infty \longrightarrow \mathbb{Z}_2 \otimes \mathcal{L}'_{n+1} \longrightarrow 0$$

and

$$0 \longrightarrow \mathbb{Z}_2 \otimes \mathcal{L}'_{n+1} \longrightarrow \tilde{\mathcal{T}}_{2n-1} \longrightarrow \mathcal{T}_{2n-1}^\infty \longrightarrow 0$$

This result is proved in [11] using the *universality* of \mathcal{T}_{2n}^∞ as the target of quadratic refinements of the canonical ‘inner product’ pairing

$$\langle \cdot, \cdot \rangle : \mathcal{L}'_{n+1} \times \mathcal{L}'_{n+1} \rightarrow \mathcal{T}_{2n}$$

given by gluing the roots of two rooted trees, see Definition 17. We can connect these short exact sequences to a single exact sequence

$$0 \longrightarrow \mathcal{T}_{2n} \longrightarrow \mathcal{T}_{2n}^\infty \longrightarrow \tilde{\mathcal{T}}_{2n-1} \longrightarrow \mathcal{T}_{2n-1}^\infty \longrightarrow 0$$

just like in Proposition 14. However, Theorem 16 makes the additional predictions that

$$\text{Cok}(\mathcal{W}_{2n} \rightarrow \mathcal{W}_{2n}^\infty) \cong \mathbb{Z}_2 \otimes \mathcal{L}'_{n+1} \cong \text{Ker}(\mathcal{W}_{2n-1} \rightarrow \mathcal{W}_{2n-1}^\infty)$$

assuming the conjectures in this paper. As a consequence of these conjectures, we would obtain new concordance invariants with values in $\mathbb{Z}_2 \otimes \mathcal{L}'_{n+1}$ and defined on \mathcal{W}_{2n}^∞ , as the obstructions for a link to bound a framed Whitney tower of order $2n$. In [11] we show that a quotient of $\mathbb{Z}_2 \otimes \mathcal{L}'_{n+1}$, namely $\mathbb{Z}_2 \otimes \mathcal{L}_{n+1}$, is indeed detected by what we call *higher-order Sato-Levine invariants*. Here the free quasi-Lie algebra \mathcal{L}' is replaced by the usual *free Lie algebra* \mathcal{L} satisfying the Jacobi identity and self-annihilation relations $[X, X] = 0$.

Levine showed in [19] that for odd n the squaring map $X \mapsto [X, X]$ induces an isomorphism

$$\mathbb{Z}_2 \otimes \mathcal{L}_{\frac{n+1}{2}} \cong \text{Ker}(\mathbb{Z}_2 \otimes \mathcal{L}'_{n+1} \rightarrow \mathbb{Z}_2 \otimes \mathcal{L}_{n+1})$$

We thus conjecture that this group is also detected by concordance invariants which are the above-mentioned *higher-order Arf invariants*. These would generalize the usual Arf invariants of the link components, which (as shown in [9]) is the case $n = 1$.

It is interesting to note that the case $n = 0$ leads to the prediction

$$\text{Cok}(\mathcal{W}_0 \rightarrow \mathcal{W}_0^\infty) \cong \mathbb{Z}_2 \otimes \mathcal{L}_1 \cong (\mathbb{Z}_2)^m$$

This is indeed the group of framed m -component links modulo those with even framings! In fact, the consistency of this computation was the motivating factor to consider filtrations of the set of *framed* links \mathbb{L} , rather than just oriented links.

As described in [11], the higher-order Sato-Levine invariants turn out to be determined by the Milnor invariants, providing a new interpretation for these classical invariants as obstructions to “untwisting” a Whitney tower. On the other hand, the higher-order Arf invariants do not appear to correspond to any known invariants. They take values in a (currently unknown) quotient of $\mathbb{Z}_2 \otimes \mathbb{L}_{\frac{n+1}{2}}$ and we show in [11] that together with the Milnor invariants they give a complete characterization of our three filtrations.

1.4. Summary of all papers in this series. Paper [8] gives an overview of the results of this paper together with the closely related papers [9, 10, 11, 12]. The classifications of the geometric filtrations of link concordance defined in the current paper are achieved in a sequence of steps:

- (i) The current paper extends our previous intersection theory of Whitney towers [7, 25] to the reduced (and twisted) settings, and proves that our obstruction theory works: If the order n intersection invariant of a (twisted) Whitney tower \mathcal{W} vanishes in $\tilde{\mathcal{T}}$ (resp. \mathcal{T}^∞) then \mathcal{W} “can be raised” from order n to order $n + 1$, without changing the link $\partial\mathcal{W}$. As a consequence, we obtain the realization epimorphisms:

$$\tilde{R}_n : \tilde{\mathcal{T}}_n \twoheadrightarrow \mathbb{W}_n \quad \text{and} \quad R_n^\infty : \mathcal{T}_n^\infty \twoheadrightarrow \mathbb{W}_n^\infty$$

We also introduce the exact sequences which explain the relationship between framed and twisted Whitney towers, motivating the definitions of the higher-order Sato-Levine and higher-order Arf invariants.

- (ii) In [9] we first show that the order n Milnor invariants for links can be thought of as an epimorphism

$$\mu_n : \mathbb{W}_n^\infty \twoheadrightarrow \mathbb{D}_n := \ker([\cdot, \cdot] : \mathbb{L}_1 \otimes \mathbb{L}_{n+1} \rightarrow \mathbb{L}_{n+2})$$

Note that \mathbb{D}_n is a free abelian group whose rank can be computed via the Hall basis algorithm. Secondly, we use the geometric notion of *gropes duality* to show that the composition

$$\eta_n : \mathcal{T}_n^\infty \xrightarrow{R_n^\infty} \mathbb{W}_n^\infty \xrightarrow{\mu_n} \mathbb{D}_n$$

can be described combinatorially in a very simple way: it is given by summing over all ways of choosing a root on a given

tree. This map η_n is closely related to the procedure for converting a Whitney tower into a grope which gives the isomorphism $W_n \cong G_n$. The construction of boundary links realizing the image of the higher-order Arf invariants leads to new geometric characterizations of links with vanishing Milnor invariants through length $2n$.

- (iii) In [10] we use combinatorial Morse theory to prove the *Levine Conjecture*, that a map η'_n (analogous to η_n), which is also defined by summing over all root choices, gives an isomorphism $T_n \cong D'_n$, where D'_n is the bracket kernel on the free quasi-Lie algebra L' (analogous to D_n). This result has implications in the study of 3-dimensional homology cylinders, where Levine originally formulated his conjecture, as well as playing a key role in providing the algebraic framework for the classifications of W_n and W_n^∞ completed in [11].
- (iv) In [11], we assemble the relevant groups into commutative diagrams of exact sequences and complete the classifications of the geometric filtrations of link concordance defined in the current paper. For instance, our resolution in [10] of the Levine Conjecture is used to show that η_n is an isomorphism for $n \equiv 0, 1, 3 \pmod{4}$ and that its kernel is $\mathbb{Z}_2 \otimes L_k$ for $n = 4k - 2$. This leads to the formulations of the higher-order Arf invariants in both the framed and twisted settings, as well as a demonstration of their equivalence, and their relation to the higher-order Sato-Levine invariants. The geometric filtrations are shown to be classified by the Milnor and higher-order Arf invariants. This allows us to give complete proofs of Theorems 6, 8 and 10 above. Moreover, we show in what sense the twisted intersection invariant τ_n^∞ is the universal quadratic refinement of its framed partner τ_n , and complete the algebraic description of the relationship between the framed and twisted Whitney tower filtrations (Theorem 16 above).
- (v) Further applications to filtrations of string links and homology cylinders are described in [12].

We emphasize that although the Milnor and higher-order Arf invariants completely classify the groups W_n and W_n^∞ , there remains the question of determining the exact (finite) 2-torsion group which is the range of the higher-order Arf invariants, as touched on above and elaborated on throughout these papers.

Acknowledgments: This paper was partially written while the first two authors were visiting the third author at the Max-Planck-Institut für Mathematik in Bonn. They all thank MPIM for its stimulating research environment and generous support. The exposition of this paper was significantly improved by a careful and insightful anonymous referee. The first author was also supported by NSF grant DMS-0604351 and the last author was also supported by NSF grants DMS-0806052 and DMS-0757312. The second author was partially supported by PSC-CUNY research grant PSCREG-41-386.

2. WHITNEY TOWERS

We sketch here the relevant theory of Whitney towers as developed in [7, 21, 25], giving details for the new notion of *twisted* Whitney towers. We work in the *smooth oriented* category (with orientations usually suppressed from notation), even though all our results hold in the locally flat topological category by the basic results on topological immersions in Freedman–Quinn [14]. In fact, it can be shown that the filtrations \mathbb{G}_n , \mathbb{W}_n and \mathbb{W}_n^ω are identical in the smooth and locally flat settings. This is because a topologically flat surface can be promoted to a smooth surface at the cost of only creating unpaired intersections of arbitrarily high order (see Remark 34).

2.1. Operations on trees. To describe Whitney towers it is convenient to use the bijective correspondence between formal non-associative bracketings of elements from the index set $\{1, 2, 3, \dots, m\}$ and rooted trees, trivalent and oriented as in Definition 2, with each univalent vertex labeled by an element from the index set, except for the *root* univalent vertex which is left unlabeled.

Definition 17. Let I and J be two rooted trees.

- (i) The *rooted product* (I, J) is the rooted tree gotten by identifying the root vertices of I and J to a single vertex v and sprouting a new rooted edge at v . This operation corresponds to the formal bracket (Figure 5 upper right). The orientation of (I, J) is inherited from those of I and J as well as the order in which they are glued.
- (ii) The *inner product* $\langle I, J \rangle$ is the unrooted tree gotten by identifying the roots of I and J to a single non-vertex point. Note that $\langle I, J \rangle$ inherits an orientation from I and J , and that all the univalent vertices of $\langle I, J \rangle$ are labeled. (Figure 5 lower right.)

- (iii) The *order* of a tree, rooted or unrooted, is defined to be the number of trivalent vertices.

The notation of this paper will not distinguish between a bracketing and its corresponding rooted tree (as opposed to the notation I and $t(I)$ used in [21, 25]). In [21, 25] the inner product is written as a dot-product, and the rooted product is denoted by $*$.




FIGURE 5. The *rooted product* (I, J) and *inner product* $\langle I, J \rangle$ of $I = (I_1, I_2)$ and $J = (J_1, J_2)$. All trivalent orientations correspond to a clockwise orientation of the plane.

2.2. Whitney disks and higher-order intersections. A collection $A_1, \dots, A_m \looparrowright (M, \partial M)$ of connected surfaces in a 4-manifold M is a *Whitney tower of order zero* if the A_i are *properly immersed* in the sense that the boundary is embedded in ∂M and the interior is generically immersed in $M \setminus \partial M$.

To each order zero surface A_i is associated the order zero rooted tree consisting of an edge with one vertex labeled by i , and to each transverse intersection $p \in A_i \cap A_j$ is associated the order zero tree $t_p := \langle i, j \rangle$ consisting of an edge with vertices labelled by i and j . Note that for singleton brackets (rooted edges) we drop the bracket from notation, writing i for (i) .

The order 1 rooted Y-tree (i, j) , with a single trivalent vertex and two univalent labels i and j , is associated to any Whitney disk $W_{(i,j)}$ pairing intersections between A_i and A_j . This rooted tree can be thought of as being embedded in M , with its trivalent vertex and rooted edge sitting in $W_{(i,j)}$, and its two other edges descending into A_i and A_j as sheet-changing paths. (The cyclic orientation at the trivalent vertex of the bracket (i, j) corresponds to an orientation of $W_{(i,j)}$ via a convention described below in 2.5.)

Recursively, the rooted tree (I, J) is associated to any Whitney disk $W_{(I,J)}$ pairing intersections between W_I and W_J (see left-hand side of Figure 6); with the understanding that if, say, I is just a singleton i ,

then W_I denotes the order zero surface A_i . Note that a $W_{(I,J)}$ can be created by a finger move pushing W_J through W_I .

To any transverse intersection $p \in W_{(I,J)} \cap W_K$ between $W_{(I,J)}$ and any W_K is associated the un-rooted tree $t_p := \langle (I, J), K \rangle$ (see right-hand side of Figure 6).




FIGURE 6. On the left, (part of) the rooted tree (I, J) associated to a Whitney disk $W_{(I,J)}$. On the right, (part of) the unrooted tree $t_p = \langle (I, J), K \rangle$ associated to an intersection $p \in W_{(I,J)} \cap W_K$. Note that p corresponds to where the roots of (I, J) and K are identified to a (non-vertex) point in $\langle (I, J), K \rangle$.

Definition 18. The *order* of a Whitney disk W_I is defined to be the order of the rooted tree I , and the *order* of a transverse intersection p is defined to be the order of the tree t_p .

Definition 19. A collection \mathcal{W} of properly immersed surfaces together with higher-order Whitney disks is an *order n Whitney tower* if \mathcal{W} contains no unpaired intersections of order less than n .

The Whitney disks in \mathcal{W} must have disjointly embedded boundaries, and generically immersed interiors. All Whitney disks and order zero surfaces must also be *framed* (as discussed next).

2.3. Twisted Whitney disks and framings. The normal disk-bundle of a Whitney disk W in M is isomorphic to $D^2 \times D^2$, and comes equipped with a canonical nowhere-vanishing *Whitney section* over the boundary given by pushing ∂W tangentially along one sheet and normally along the other, avoiding the tangential direction of W (see Figure 7, and e.g. 1.7 of [26]). Pulling back the orientation of M with the requirement that the normal disks have +1 intersection with W means the Whitney section determines a well-defined (independent of the orientation of W) relative Euler number $\omega(W) \in \mathbb{Z}$ which represents the obstruction to extending the Whitney section across W .

Following traditional terminology, when $\omega(W)$ vanishes W is said to be *framed*. (Since $D^2 \times D^2$ has a unique trivialization up to homotopy, this terminology is only mildly abusive.) In general when $\omega(W) = k$, we say that W is *k -twisted*, or just *twisted* if the value of $\omega(W)$ is not specified. So a 0-twisted Whitney disks is a framed Whitney disk.




FIGURE 7. The Whitney section over the boundary of a framed Whitney disk is indicated by the dotted loop shown on the left for a clean Whitney disk W in a 3-dimensional slice of 4-space. On the right is shown an embedding into 3-space of the normal disk-bundle over ∂W , indicating how the Whitney section determines a well-defined nowhere vanishing section which lies in the I -sheet and is normal to the J -sheet.

Note that a *framing* of ∂A_i (respectively A_i) is by definition a trivialization of the normal bundle of the immersion. If the ambient 4-manifold is oriented, this is equivalent to an orientation and a non-vanishing normal vector field on ∂A_i (respectively A_i). The twisting $\omega(A_i) \in \mathbb{Z}$ of an order zero surface is also defined when a framing of ∂A_i is given, and A_i is said to be *framed* when $\omega(A_i) = 0$.

2.4. Twisted Whitney towers. In the definition of an order n Whitney tower given just above (following [7, 21, 22, 25]) all Whitney disks and order zero surfaces are required to be framed. It turns out that the natural generalization to twisted Whitney towers involves allowing twisted Whitney disks only in at least “half the order” as follows:

Definition 20. A *twisted Whitney tower of order 0* is a collection of properly immersed surfaces in a 4-manifold (without any framing requirement).

For $n > 0$, a *twisted Whitney tower of order $(2n-1)$* is just a (framed) Whitney tower of order $(2n-1)$ as in Definition 19 above.

For $n > 0$, a *twisted Whitney tower of order $2n$* is a Whitney tower having all intersections of order less than $2n$ paired by Whitney disks,

with all Whitney disks of order less than n required to be framed, but Whitney disks of order at least n allowed to be twisted.

Remark 21. Note that, for any n , an order n (framed) Whitney tower is also an order n twisted Whitney tower. We may sometimes refer to a Whitney tower as a *framed* Whitney tower to emphasize the distinction, and will always use the adjective “twisted” in the setting of Definition 20.

Remark 22. The convention of allowing only order $\geq n$ twisted Whitney disks in order $2n$ twisted Whitney towers is explained both algebraically and geometrically in [9]. In any event, an order $2n$ twisted Whitney tower can always be modified so that all its Whitney disks of order $> n$ are framed, so the twisted Whitney disks of order equal to n are the important ones.

2.5. Whitney tower orientations. Orientations on order zero surfaces in a Whitney tower \mathcal{W} are fixed, and required to induce the orientations on their boundaries. After choosing and fixing orientations on all the Whitney disks in \mathcal{W} , the associated trees are embedded in \mathcal{W} so that the vertex orientations are induced from the Whitney disk orientations, with the descending edges of each trivalent vertex enclosing the *negative intersection point* of the corresponding Whitney disk, as in Figure 6. (In fact, if a tree t has more than one trivalent vertex which corresponds to the same Whitney disk, then t will only be immersed in \mathcal{W} , but this immersion can be taken to be a local embedding around each trivalent vertex of t as in Figure 6.)

This “negative corner” convention, which differs from the positive corner convention in [7, 25], will turn out to be compatible with commutator conventions for use in [9].

With these conventions, different choices of orientations on Whitney disks in \mathcal{W} correspond to anti-symmetry relations (as explained in [25]).

2.6. Intersection invariants for Whitney towers. We recall from Definition 2 that the abelian group \mathcal{T}_n is the free abelian group on labeled vertex-oriented order n trees, modulo the AS and IHX relations, see Figure 2. The obstruction theory of [25] in the current simply connected setting works as follows.

Definition 23. The *order n intersection invariant* $\tau_n(\mathcal{W})$ of an order n Whitney tower \mathcal{W} is defined to be

$$\tau_n(\mathcal{W}) := \sum \epsilon_p \cdot t_p \in \mathcal{T}_n$$

where the sum is over all order n intersections p , with $\epsilon_p = \pm 1$ the usual sign of a transverse intersection point.

As stated in Theorem 4 in the introduction, if L bounds $\mathcal{W} \subset B^4$ with $\tau_n(\mathcal{W}) = 0 \in \mathcal{T}_n$, then L bounds a Whitney tower of order $n + 1$. This is a special case of the simply connected version of the more general Theorem 2 of [25]. We will use the following version of Theorem 2 of [25] where the order zero surfaces are either properly immersed disks in B^4 or properly immersed annuli in $S^3 \times I$:

Theorem 24 ([25]). *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports an order n Whitney tower \mathcal{W} with $\tau_n(\mathcal{W}) = 0 \in \mathcal{T}_n$, then A is regularly homotopic (rel ∂) to A' which supports an order $n + 1$ Whitney tower.*

2.7. Intersection invariants for twisted Whitney towers. The intersection invariants for Whitney towers are extended to twisted Whitney towers as follows:

Definition 25. The abelian group $\mathcal{T}_{2n-1}^\infty$ is the quotient of \mathcal{T}_{2n-1} by the *boundary-twist relations*:

$$\langle (i, J), J \rangle = i \multimap_J^J = 0$$

Here J ranges over all order $n - 1$ rooted trees. (This is the same as taking the quotient of $\tilde{\mathcal{T}}_{2n-1} = \mathcal{T}_{2n-1}/\text{Im}(\Delta_{2n-1})$ by boundary-twist relations since $\text{Im}(\Delta_{2n-1})$ is contained in the span of boundary-twist relations – see Section 5).

The boundary-twist relations correspond geometrically to the fact that performing a boundary twist (Figure 13) on an order n Whitney disk $W_{(i,J)}$ creates an order $2n - 1$ intersection point $p \in W_{(i,J)} \cap W_J$ with associated tree $t_p = \langle (i, J), J \rangle$ (which is 2-torsion by the AS relations) and changes $\omega(W_{(i,J)})$ by ± 1 . Since order n twisted Whitney disks are allowed in an order $2n$ Whitney tower such trees do not represent obstructions to the existence of the next order twisted tower.

For any rooted tree J we define the corresponding *∞ -tree*, denoted by J^∞ , by labeling the root univalent vertex with the symbol “ ∞ ”:

$$J^\infty := \infty — J$$

Definition 26. The abelian group \mathcal{T}_{2n}^∞ is the free abelian group on order $2n$ trees and order n ∞ -trees, modulo the following relations:

- (i) AS and IHX relations on order $2n$ trees (Figure 2)
- (ii) *symmetry* relations: $(-J)^\infty = J^\infty$
- (iii) *twisted IHX* relations: $I^\infty = H^\infty + X^\infty - \langle H, X \rangle$
- (iv) *interior twist* relations: $2 \cdot J^\infty = \langle J, J \rangle$

Here the AS and IHX relations are as usual, but they only apply to non- ∞ trees. The *symmetry relation* corresponds to the fact that

the relative Euler number $\omega(W)$ is independent of the orientation of the Whitney disk W , with $-J$ denoting the “opposite” orientation of J (meaning that the trivalent orientations differ at an odd number of vertices). The *twisted IHX relation* corresponds to the effect of performing a Whitney move in the presence of a twisted Whitney disk, as described below in Lemma 42. The *interior-twist relation* corresponds to the fact that creating a ± 1 self-intersection in a W_J changes the twisting by ∓ 2 (Figure 14).

Remark 27. The symmetry, twisted IHX, and interior twist relations in \mathcal{T}_{2n}^∞ have a surprisingly natural algebraic interpretation that we explain in [11]. The idea is to extend the map $J \mapsto J^\infty$ to a *symmetric quadratic refinement* q of the bilinear form $\langle \cdot, \cdot \rangle$ on the free quasi-Lie algebra of rooted trees (the intersection form on Whitney disks) by defining $q(J) = J^\infty$ and extending to linear combinations by the formula

$$q(J + K) := J^\infty + K^\infty + \langle J, K \rangle$$

Expanding $q(I - H + X) = 0$ leads to the 6-term IHX relation

$$I^\infty + H^\infty + X^\infty = \langle I, H \rangle - \langle I, X \rangle + \langle H, X \rangle$$

which is equivalent to the twisted IHX relation in the presence of the interior-twist relations. Those in turn follow by setting $K := -J$ from the symmetry relation.

Remark 28. We discovered in [7] that the (framed) IHX relation can be realized in three dimensions as well as four, and it is interesting to note that many of the relations that we obtain for twisted Whitney towers in four dimensions can also be realized by rooted clasper surgeries (grope cobordisms) in three dimensions. Here the twisted Whitney disk corresponds to a ± 1 framed leaf of a clasper. For example the relation $I^\infty = H^\infty + X^\infty - \langle H, X \rangle$ has the following clasper explanation. I^∞ represents a clasper with one isolated twisted leaf. By the topological IHX relation, one can replace I^∞ by two clasps of the form H^∞ and $(-X)^\infty = X^\infty$ embedded in a regular neighborhood of the original clasper with leaves parallel to the leaves of the original. The twisted leaves are now linked together, so applying Habiro’s zip construction (which complicates the picture considerably) one gets three tree clasps, of the form H^∞ , X^∞ and $\langle H, -X \rangle$ respectively.

Similarly, the relation $2 \cdot J^\infty = \langle J, J \rangle$ has an interpretation where one takes a clasper which represents J^∞ and splits off a geometrically cancelling parallel copy, representing the tree J^∞ . Again, because the twisted leaves link, we also get the term $\langle J, -J \rangle$.

These observations will be enlarged upon in [12] to analyze filtrations on homology cylinders.

Recall from Definition 20 (and Remark 22) that twisted Whitney disks only occur in even order twisted Whitney towers, and only those of half-order are relevant to the obstruction theory.

Definition 29. The *order n intersection invariant* $\tau_n^\infty(\mathcal{W})$ of an order n twisted Whitney tower \mathcal{W} is defined to be

$$\tau_n^\infty(\mathcal{W}) := \sum \epsilon_p \cdot t_p + \sum \omega(W_J) \cdot J^\infty \in \mathcal{T}_n^\infty$$

where the first sum is over all order n intersections p and the second sum is over all order $n/2$ Whitney disks W_J with twisting $\omega(W_J) \in \mathbb{Z}$. For $n = 0$, recall from 2.2 above our notational convention that W_j denotes A_j , and that $\omega(A_j) \in \mathbb{Z}$ is the relative Euler number of the normal bundle of A_j with respect to the given framing of ∂A_j as in 2.3 .

By splitting the twisted Whitney disks, as explained in subsection 2.9 below, for $n > 0$ we may actually assume that all non-zero $\omega(W_J) \in \{\pm 1\}$, just like the signs ϵ_p .

As in the framed case, the vanishing of τ_n^∞ is sufficient for the existence of a twisted Whitney tower of order $(n + 1)$, and the proof of Theorem 10 in Section 4 will be based on the following analogue of the framed order-raising Theorem 24 to the twisted setting:

Theorem 30. *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports an order n twisted Whitney tower \mathcal{W} with $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$, then A is regularly homotopic (rel ∂) to A' which supports an order $n + 1$ twisted Whitney tower.*

The proof of Theorem 30 is given in Section 4 below.

Proofs of the “order-raising” Theorems 30 and 24 (and its strengthening Theorem 45 below) depend on realizing the relations in the target groups by controlled manipulations of Whitney towers. The next two subsections introduce combinatorial notions useful for describing the algebraic effect of such geometric constructions.

For the rest of this section we assume our Whitney towers are of positive order for convenience of notation.

2.8. Intersection forests. Recall that the trees associated to intersections and Whitney disks in a Whitney tower can be considered to be immersed in the Whitney tower, with vertex orientations induced by the Whitney tower orientation, as in Figure 6.

Definition 31. The *intersection forest* $t(\mathcal{W})$ of a framed Whitney tower \mathcal{W} is the disjoint union of signed trees associated to all unpaired intersections p in \mathcal{W} :

$$t(\mathcal{W}) = \coprod_p \epsilon_p \cdot t_p$$

with ϵ_p the sign of the intersection point p . For \mathcal{W} of order n , we can think of the signed order n trees in $t(\mathcal{W})$ as an “abelian word” in the generators $\pm t_p$ which represents $\tau_n(\mathcal{W}) \in \mathcal{T}_n$. More precisely, $t(\mathcal{W})$ is an element of the free abelian monoid, with unit \emptyset , generated by (isomorphism classes of) signed trees, trivalent, labeled and vertex-oriented as usual. We emphasize that there are no cancellations or other relations here.

Remark 32. In the older papers [7, 21, 25] we referred to $t(\mathcal{W})$ as the “geometric intersection tree” (and to the group element $\tau_n(\mathcal{W})$ as the order n intersection “tree”, rather than “invariant”), but the term “forest” better describes the disjoint union of (signed) trees $t(\mathcal{W})$.

Similarly to the framed case, the *intersection forest* $t(\mathcal{W})$ of a *twisted* Whitney tower \mathcal{W} is the disjoint union of signed trees associated to all unpaired intersections p in \mathcal{W} and integer-coefficient ∞ -trees associated to all non-trivially twisted Whitney disks W_J in \mathcal{W} :

$$t(\mathcal{W}) = \coprod_p \epsilon_p \cdot t_p + \coprod_J \omega(W_J) \cdot J^\infty$$

with $\omega(W_J) \in \mathbb{Z}$ the twisting of W_J . Again, there are no cancellations or relations (and the informal “+” sign in the expression is purely cosmetic).

We will see in the next subsection that all the trees can be made to be disjoint in \mathcal{W} , with all non-zero $\omega(W_J) = \pm 1$, so that $t(\mathcal{W})$ is also a topological disjoint union which corresponds to an element in the free abelian monoid generated by (isomorphism classes of) signed trees, and signed ∞ -trees.

2.9. Splitting twisted Whitney towers. A framed Whitney tower is *split* if the set of singularities in the interior of any Whitney disk consists of either a single point, or a single boundary arc of a Whitney disk, or is empty. This can always be arranged, as observed in Lemma 13 of [25] (Lemma 3.5 of [21]), by performing finger moves along Whitney disks guided by arcs connecting the Whitney disk boundary arcs. Implicit in this construction is that the finger moves preserve the Whitney disk framings (by not twisting relative to the Whitney disk that is being split – see Figure 8). A Whitney disk W is *clean* if the interior of W is embedded and disjoint from the rest of the Whitney tower. In the setting of twisted Whitney towers, it will simplify the combinatorics to




FIGURE 8. A neighborhood of a twisted finger move which splits a Whitney disk into two Whitney disks. The vertical black arcs are slices of the new Whitney disks, and the grey arcs are slices of extensions of the Whitney sections. The finger-move is supported in a neighborhood of an arc in the original Whitney disk running from a point in the Whitney disk boundary on the “upper” surface sheet to a point in the Whitney disk boundary on the “lower” surface sheet. (Before the finger-move this guiding arc would have been visible in the middle picture as a vertical black arc-slice of the original Whitney disk.)

use “twisted” finger moves to similarly split-off twisted Whitney disks into ± 1 -twisted clean Whitney disks.

We call a twisted Whitney tower *split* if all of its non-trivially twisted Whitney disks are clean and have twisting ± 1 , and all of its framed Whitney disks are split in the usual sense (as for framed Whitney towers).

Lemma 33. *If A supports an order n twisted Whitney tower \mathcal{W} , then A is homotopic (rel ∂) to A' which supports a split order n twisted Whitney tower \mathcal{W}' , such that:*

- (i) *The disjoint union of non- ∞ trees $\coprod_p \epsilon_p \cdot t_p \subset t(\mathcal{W})$ is isomorphic to the disjoint union of non- ∞ trees $\coprod_{p'} \epsilon_{p'} \cdot t_{p'} \subset t(\mathcal{W}')$.*
- (ii) *Each $\omega(W_J) \cdot J^\infty$ in $t(\mathcal{W})$ gives rise to the disjoint union of exactly $|\omega(W_J)|$ -many $\pm 1 \cdot J^\infty$ in \mathcal{W}' , where the sign \pm corresponds to the sign of $\omega(W_J)$.*

Proof. Illustrated in Figure 8 is a local picture of a twisted finger move, which splits one Whitney disk into two, while also changing twistings. If the original Whitney disk in Figure 8 was framed, then the two new Whitney disks will have twistings $+1$ and -1 , respectively. In general, if the arc guiding the finger move splits the twisting of the original Whitney disk into ω_1 and ω_2 zeros of the extended Whitney

section, then the two new Whitney disks will have twistings $\omega_1 + 1$ and $\omega_2 - 1$, respectively. Thus, by repeatedly splitting off framed corners into ± 1 -twisted Whitney disks, any ω -twisted Whitney disk ($\omega \in \mathbb{Z}$) can be split into $|\omega|$ -many $+1$ -twisted or -1 -twisted clean Whitney disks, together with split framed Whitney disks containing any interior intersections in the original twisted Whitney disk. Combining this with the untwisted splitting of the framed Whitney disks as in Lemma 13 of [25] gives the result. \square

Remark 34. We sketch here a brief explanation of why the smooth and locally flat filtrations are equal. A locally flat surface can be made smooth by a small perturbation, which after introducing cusps as necessary can be assumed to be a regular (locally flat) homotopy. By a general position argument, this regular homotopy can be assumed to be a finite number of finger moves, which are guided by arcs and lead to canceling self-intersection pairs which admit small disjointly embedded Whitney disks (which are ‘inverses’ to the finger moves). These Whitney disks are only locally flat, but can be perturbed to be smooth, again only at the cost of creating paired self-intersections, and iteration of this process leads to an arbitrarily high-order smooth sub-Whitney tower pairing all intersections created by the original surface perturbation.

3. THE REALIZATION MAPS

This section contains clarifications and proofs of Theorems 3 and 10 from the introduction which state the existence of surjections $R_n: \mathcal{T}_n \rightarrow W_n$ and $R_n^\infty: \mathcal{T}_n^\infty \rightarrow W_n^\infty$ for all n , in particular exhibiting the sets W_n and W_n^∞ as finitely generated abelian groups under connected sum.

All proofs in this section apply in the reduced setting as well, and the constructions described here also define the surjections $\tilde{R}_n: \tilde{\mathcal{T}}_n \rightarrow W_n$ described in the introduction.

Recall that our manifolds are assumed oriented, but orientations are suppressed from the discussion as much as possible. In the following an orientation is fixed once and for all on S^3 ; and a *framed link* has oriented components, each equipped with a nowhere-vanishing normal section.

Definition 35. A framed link $L \subset S^3 = \partial B^4$ bounds an order n Whitney tower \mathcal{W} if $\mathcal{W} \subset B^4$ is an order n Whitney tower whose order zero surfaces are immersed disks bounded by the components of L , as in Definition 19.

Similarly, a framed link $L \subset S^3 = \partial B^4$ bounds an order n twisted Whitney tower \mathcal{W} if $\mathcal{W} \subset B^4$ is an order n twisted Whitney tower

whose order zero surfaces are immersed disks bounded by the components of L , as in Definition 20.

Definition 36. For $n \geq 1$, framed links L_0 and L_1 in S^3 are *Whitney tower concordant of order n* if the i th components of $L_0 \subset S^3 \times \{0\}$ and $-L_1 \subset S^3 \times \{1\}$ cobound an immersed annulus A_i for each i such that the A_i are transverse and support an order n Whitney tower. If the A_i support a *twisted* order n Whitney tower then L_0 and L_1 are said to be *twisted Whitney tower concordant of order n* .

Note that a (twisted) Whitney tower concordance preserves framings on L_0 and L_1 (as links in S^3) since all self-intersections of the A_i are paired by Whitney disks in any (twisted) Whitney tower of order $n \geq 1$.

Recall from the introduction that the set of m -component framed links in S^3 which bound order n (twisted) Whitney towers in B^4 is denoted by $\mathbb{W}_n = \mathbb{W}_n(m)$ (resp. \mathbb{W}_n^ω); and the quotient of \mathbb{W}_n by the equivalence relation of order $n+1$ (twisted) Whitney tower concordance is denoted by \mathcal{W}_n (resp. \mathcal{W}_n^ω).

Throughout this section the twisted setting mirrors the framed setting, with discussions and arguments given simultaneously.

We first need to show our essential criterion, Corollary 5, for links to represent equal elements in the associated graded \mathcal{W}_n : Links L_0 and L_1 in \mathbb{W}_n represent the same element of \mathcal{W}_n if and only if there exist order n Whitney towers \mathcal{W}_i in B^4 with $\partial \mathcal{W}_i = L_i$ and $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1) \in \mathcal{T}_n$.

Proof of Corollary 5. If L_0 and L_1 are equal in \mathcal{W}_n then they cobound A supporting an order $n+1$ Whitney tower \mathcal{V} in $S^3 \times I$, and any order n Whitney tower \mathcal{W}_1 in B^4 bounded by L_1 can be extended by \mathcal{V} to form an order n Whitney tower \mathcal{W}_0 in B^4 bounded by L_0 , with $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1) \in \mathcal{T}_n$ since $\tau_n(\mathcal{V})$ vanishes.

Conversely, suppose that L_0 and L_1 bound order n Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in 4-balls B_0^4 and B_1^4 , with $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1)$. Then constructing $S^3 \times I$ as the connected sum $B_0^4 \# B_1^4$ (along balls in the complements of \mathcal{W}_0 and \mathcal{W}_1), and tubing together the corresponding order zero disks of \mathcal{W}_0 and \mathcal{W}_1 , and taking the union of the Whitney disks in \mathcal{W}_0 and \mathcal{W}_1 , yields a collection A of properly immersed annuli connecting L_0 and L_1 and supporting an order n Whitney tower \mathcal{V} . Since the orientation of the ambient 4-manifold has been reversed for one of the original Whitney towers, say \mathcal{W}_1 , which results in a global sign change for $\tau_n(\mathcal{W}_1)$, it follows that \mathcal{V} has vanishing order n intersection invariant:

$$\tau_n(\mathcal{V}) = \tau_n(\mathcal{W}_0) - \tau_n(\mathcal{W}_1) = \tau_n(\mathcal{W}_0) - \tau_n(\mathcal{W}_0) = 0 \in \mathcal{T}_n$$

So by Theorem 24, A is homotopic (rel ∂) to A' supporting an order $n + 1$ Whitney tower, and hence L_0 and L_1 are equal in \mathbb{W}_n . \square

Remark 37. The analogous statement and proof of Corollary 5 holds in the twisted case (with Theorem 30 playing the role of Theorem 24). For this case, we'll spell out the statement carefully but in several instances below we will just state that the twisted case is analogous: Links L_0 and L_1 in \mathbb{W}_n^∞ represent the same element of \mathbb{W}_n^∞ if and only if there exist order n twisted Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in B^4 bounded by L_0 and L_1 respectively such that $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$.

Remark 38. Remark 37 similarly applies to the reduced setting by Theorem 45 below, although we will omit further reference to $\tilde{\mathcal{T}}$ in this section.

3.1. Band sums of links. The *band sum* $L \#_\beta L' \subset S^3$ of oriented m -component links L and L' along bands β is defined as follows: Form S^3 as the connected sum of 3-spheres containing L and L' along balls in the link complements. Let β be a collection of disjointly embedded oriented bands joining like-indexed link components such that the band orientations are compatible with the link orientations. Take the usual connected sum of each pair of components along the corresponding band. Although it is well-known that the concordance class of $L \#_\beta L'$ depends in general on β , it turns out that the image of $L \#_\beta L'$ in \mathbb{W}_n (or in \mathbb{W}_n^∞) does not depend on β :

Lemma 39. *For links L and L' representing elements of \mathbb{W}_n , any band sum $L \#_\beta L'$ represents an element of \mathbb{W}_n which only depends on the equivalence classes of L and L' in \mathbb{W}_n . The same statement holds in \mathbb{W}_n^∞ .*

Proof. We shall only give the proof in the framed case, the twisted case is analogous. If L_0 and L_1 represent the same element of \mathbb{W}_n , and if L'_0 and L'_1 represent the same element of \mathbb{W}_n , then by Corollary 5 above, for $i = 0, 1$, there are order n Whitney towers \mathcal{W}_i and \mathcal{W}'_i bounding L_i and L'_i such that $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1)$ and $\tau_n(\mathcal{W}'_0) = \tau_n(\mathcal{W}'_1)$. By Lemma 40 just below, $L_i \#_{\beta_i} L'_i$ bounds $\mathcal{W}_i^\#$ for $i = 0, 1$, with

$$\tau_n(\mathcal{W}_0^\#) = \tau_n(\mathcal{W}_0) + \tau_n(\mathcal{W}'_0) = \tau_n(\mathcal{W}_1) + \tau_n(\mathcal{W}'_1) = \tau_n(\mathcal{W}_1^\#)$$

so again by Corollary 5, $L_0 \#_{\beta_0} L'_0$ is order $n + 1$ Whitney tower concordant to $L_1 \#_{\beta_1} L'_1$, hence $L_0 \#_{\beta_0} L'_0$ and $L_1 \#_{\beta_1} L'_1$ represent the same element of \mathbb{W}_n . \square

Lemma 40. *If L and L' bound order n (twisted) Whitney towers \mathcal{W} and \mathcal{W}' in B^4 , then for any β there exists an order n (twisted) Whitney*

tower $\mathcal{W}^\# \subset B^4$ bounded by $L\#_\beta L'$, such that $t(\mathcal{W}^\#) = t(\mathcal{W}) \amalg t(\mathcal{W}')$, where $t(\mathcal{V})$ denotes the intersection forest of a Whitney tower \mathcal{V} as above in subsection 2.8.

Proof. Let B and B' be the 3-balls in the link complements used to form the S^3 containing $L\#_\beta L'$. Then gluing together the two 4-balls containing \mathcal{W} and \mathcal{W}' along B and B' forms B^4 containing $L\#_\beta L'$ in its boundary. Take $\mathcal{W}^\#$ to be the boundary band sum of \mathcal{W} and \mathcal{W}' along the order zero disks guided by the bands β , with the interiors of the bands perturbed slightly into the interior of B^4 . It is clear that $t(\mathcal{W}^\#)$ is just the disjoint union $t(\mathcal{W}) \amalg t(\mathcal{W}')$ since no new singularities have been created. \square

3.2. The realization maps. The realization maps R_n are defined as follows: Given any group element $g \in \mathcal{T}_n$, by Lemma 41 just below there exists an m -component link $L \subset S^3$ bounding an order n Whitney tower $\mathcal{W} \subset B^4$ such that $\tau_n(\mathcal{W}) = g \in \mathcal{T}_n$. Define $R_n(g)$ to be the class determined by L in \mathcal{W}_n . This is well-defined (does not depend on the choice of such L) by Corollary 5. The twisted realization map R_n^ω is defined the same way using twisted Whitney towers.

Lemma 41. *For any disjoint union $\amalg_p \epsilon_p \cdot t_p + \amalg_J \omega(W_J) \cdot J^\omega$ there exists an m -component link L bounding a twisted Whitney tower \mathcal{W} with intersection forest $t(\mathcal{W}) = \amalg_p \epsilon_p \cdot t_p + \amalg_J \omega(W_J) \cdot J^\omega$. If the disjoint union contains no ω -trees then all Whitney disks in \mathcal{W} are framed.*

Note that if in the disjoint union all non- ω trees are order at least n and all ω -trees are order at least $n/2$ then \mathcal{W} will have order n .

Proof. It suffices to consider the cases where the disjoint union consists of just a single (signed) tree or ω -tree since by Lemma 40 any sum of such trees can then be realized by band sums of links.

The following algorithm, in the untwisted case, is the algorithm called "Bing-doubling along a tree" by Cochran and used in Section 7 of [2] and Theorem 3.3 of [3] to produce links in S^3 with prescribed (first non-vanishing) Milnor invariants.

Realizing order zero trees and ω -trees. A 0-framed Hopf link bounds an order zero Whitney tower $\mathcal{W} = D_1 \cup D_2 \subset B^4$, where the two embedded disks D_1 and D_2 have a single interior intersection point p with $t_p = \langle 1, 2 \rangle = 1 - 2$ (see Figure 9). Assuming appropriate fixed orientations of B^4 and S^3 , the sign ϵ_p associated to p is the usual sign of the Hopf link. So taking a 0-framed $(m-2)$ -component trivial link together with a Hopf link (as the i th and j th components) gives an




FIGURE 9. Pushing into B^4 from left to right: A Hopf link in $S^3 = \partial B^4$ bounds embedded disks $D_1 \cup D_2 \subset B^4$ which intersect in a point p , with $t_p = \langle 1, 2 \rangle$.




FIGURE 10. Pushing into B^4 from left to right: The disks D_{i_2} and D_j extend to the right-most picture where they are completed by capping off the unlink. The disk D_{i_1} only extends to the middle picture where the intersections between D_{i_1} and D_{i_2} are paired by the Whitney disk $W_{(i_1, i_2)}$, that has a single interior intersection $p \in W_{(i_1, i_2)} \cap D_j$ with $t_p = \langle (i_1, i_2), j \rangle$.

m -component link L bounding \mathcal{W} with $t(\mathcal{W}) = \epsilon_p \cdot \langle i, j \rangle = \epsilon_p \cdot i — j$, for any $\epsilon_p = \pm 1$, and $i \neq j$.

To realize the tree $\pm i — i$, we can use the unlink with framings 0, except that the component labeled by the index i has framing ± 2 . Similarly, if the component has framing ± 1 then the resulting tree is $\pm \infty — i$.

Realizing order 1 trees. Consider now a link L whose i th and j th components form a Hopf link $L^i \cup L^j$ bounding disks $D_i \cup D_j \subset B^4$ with transverse intersection $p = D_i \cap D_j$. Assume that $D_i \cup D_j$ extends to an order zero Whitney tower \mathcal{W} bounded by L with $t(\mathcal{W}) = \epsilon_p \cdot t_p = \epsilon_p \cdot \langle i, j \rangle$.

Replacing L^i by an untwisted Bing-double $L^{i_1} \cup L^{i_2}$ results in a new sublink of Borromean rings $L^{i_1} \cup L^{i_2} \cup L^j$ bounding disks $D_{i_1} \cup D_{i_2} \cup D_j$ as indicated in Figure 10, with D_{i_1} and D_{i_2} intersecting in a canceling




FIGURE 11. Pushing into B^4 from left to right: An i - and j -labeled n -twisted Bing-double (case $n = 2$) of the unknot in $S^3 = \partial B^4$ bounds disks D_i and D_j whose intersections are paired by a Whitney disk $W_{(i,j)}$. D_j extends to the right-hand picture but D_i only extends to the middle picture, where the boundary of $W_{(i,j)}$ is indicated by the dark arcs. The rest of $W_{(i,j)}$ extends into the right-hand picture where disjointly embedded disks bounded by the unlink complete both $W_{(i,j)}$ and D_j . The interior of $W_{(i,j)}$ is embedded and disjoint from both D_i and D_j . Figure 12 shows that $W_{(i,j)}$ is twisted, with $\omega(W_{(i,j)}) = n$.

pair of intersections paired by an order 1 Whitney disk $W_{(i_1, i_2)}$, which can be formed from D_i with a small collar removed, so that $W_{(i_1, i_2)}$ has a single intersection with D_j corresponding to the original $p = D_i \cap D_j$. (One can think of D_{i_1} and D_{i_2} as being formed by the trace of the obvious pulling-apart homotopy that shrinks L^{i_1} and L^{i_2} down in a tubular neighborhood of L^i , with the canceling pair of intersections between D_{i_1} and D_{i_2} being created as the clasps are pulled apart.)

The effect of this Bing-doubling operation on the intersection forest is that the original order zero $t_p = \langle i, j \rangle$ has given rise to the order 1 tree $\langle (i_1, i_2), j \rangle$. Switching the orientation on one of the new components changes the sign of p , as can be checked using our orientation conventions. By relabeling and/or banding together components of this new link any labels on this order 1 tree can be realized. Since the doubling was untwisted, $W_{(i_1, i_2)}$ is framed (see Figures 11 and 12), so the Whitney tower bounded by the new link is order 1.

Realizing order n trees. Since any order n tree can be gotten from some order $n - 1$ tree by attaching two new edges to a univalent vertex as in the previous paragraph, it follows inductively that any order n tree is the intersection forest of a Whitney tower bounded by some link. (First create a distinctly-labeled tree of the desired ‘shape’ by doubling, then correct the labels by interior band-summing.)




FIGURE 12. The Whitney section over $\partial W_{(i,j)}$ (from Figure 11) is indicated by the dashed arcs on the left. The twisting $\omega(W_{(i,j)}) = n$ (the obstruction to extending the Whitney section across the Whitney disk) corresponds to the n -twisting of the Bing-doubling operation.

Realizing ∞ -trees of order 1. As illustrated (for the case $n = 2$) in Figures 11 and 12, the n -twisted Bing-double of the unknot (with components labeled i and j) bounds an order 2 twisted Whitney tower \mathcal{W} with $t(\mathcal{W}) = n \cdot (i, j)^\infty = n \cdot \infty \prec_j^i$. Banding together the two components would yield a knot realizing $(i, i)^\infty$.

Realizing ∞ -trees of order n . By applying iterated untwisted Bing-doubling operations to the i - and j -labeled components of the order 1 case, one can construct for any rooted tree (I, J) a link bounding a twisted Whitney tower \mathcal{W} with $t(\mathcal{W}) = n \cdot (I, J)^\infty$. For instance, if in the construction of Figure 11 the j -labeled link component is replaced by an untwisted Bing-double, then the disk D_j in that construction would be replaced by a (framed) Whitney disk $W_{(j_1, j_2)}$, and the n -twisted $W_{(i,j)}$ would be replaced by an n -twisted $W_{(i, (j_1, j_2))}$. (As for non- ∞ trees above, first create a distinctly-labeled tree of the desired ‘shape’ by doubling, then correct the labels by interior band-summing.) \square

3.3. Proofs of Theorem 3 and Theorem 10. Recall the content of Theorem 3: The realization maps $R_n : \mathcal{T}_n \rightarrow \mathbb{W}_n$ are epimorphisms, with the group operation on \mathbb{W}_n given by band sum $\#$. Moreover, \mathbb{W}_n is the set of framed links $L \in \mathbb{W}_n$ modulo the relation that $[L_1] = [L_2] \in \mathbb{W}_n$ if and only if $L_1\# - L_2$ lies in \mathbb{W}_{n+1} , for some choice of connected sum $\#$, where $-L$ is the mirror image of L with reversed framing. The content of Theorem 10 in the twisted setting is analogous.

From Lemma 39 the band sum of links gives well-defined operations in \mathbb{W}_n and \mathbb{W}_n^∞ which are clearly associative and commutative, with the m -component unlink representing an identity element. The realization maps are homomorphisms by Lemma 40 and surjectivity is proven as follows: Given any link $L \in \mathbb{W}_n$, choose a Whitney tower \mathcal{W} of order n with boundary L and compute $\tau := \tau_n(\mathcal{W})$. Then take $L' := R_n(\tau)$,

a link that's obviously in the image of R_n and for which we know a Whitney tower \mathcal{W}' with boundary L' and $\tau(\mathcal{W}') = \tau$. By Corollary 5 it follows that L and L' represent the same element in \mathbb{W}_n .

Considering the second “Moreover...” statements of Theorem 3 and Theorem 10, first assume that L_0 and L_1 represent the same element of \mathbb{W}_n (resp. \mathbb{W}_n^∞). Then by Corollary 5, there exist order n (twisted) Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in B^4 bounded by L_0 and L_1 respectively such that $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1) \in \mathcal{T}_n$ (resp. $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$). We want to show that $L_0 \# -L_1$ bounds an order $n+1$ (twisted) Whitney tower, which will follow from Lemma 40 and the “order-raising” Theorem 24 (respectively Theorem 30) if $-L_1$ bounds an order n (twisted) Whitney tower $\overline{\mathcal{W}_1}$ such that $\tau_n(\overline{\mathcal{W}_1}) = -\tau_n(\mathcal{W}_1) \in \mathcal{T}_n$ (resp. $\tau_n^\infty(\overline{\mathcal{W}_1}) = -\tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$). If r denotes the reflection on S^3 which sends L_1 to $-L_1$, then the product $r \times \text{id}$ of r with the identity is an involution on $S^3 \times I$, and the image $r \times \text{id}(\mathcal{W}_1)$ of \mathcal{W}_1 is such a $\overline{\mathcal{W}_1}$. To see this, observe that $r \times \text{id}$ switches the signs of all transverse intersection points, and is an isomorphism on the oriented trees in \mathcal{W}_1 ; and hence switches the signs of all Whitney disk framing obstructions (which can be computed as intersection numbers between Whitney disks and their push-offs) – note that $r \times \text{id}$ is only being applied to \mathcal{W}_1 , while $S^3 \times I$ is fixed.

For the other direction of the “Moreover...” statements, assume that $L_0 \# -L_1 \subset S^3$ bounds an order $n+1$ (twisted) Whitney tower $\mathcal{W} \subset B^4$. By the definition of connected sum, S^3 decomposes as the union of two disjoint 3-balls B_0 and B_1 containing L_0 and $-L_1$, joined together by the $S^2 \times I$ through which passes the bands guiding the connected sum. Taking another 4-ball with the same decomposition of its boundary 3-sphere, and gluing the 4-balls together by identifying the boundary 2-spheres of the 3-balls, and identifying the $S^2 \times I$ subsets by the identity map, forms $S^3 \times I$ containing an order $n+1$ (twisted) Whitney tower concordance between L_0 and $-L_1$ which consists of \mathcal{W} together with the parts of the connected-sum bands that are contained in $S^2 \times I$.

4. PROOF OF THEOREM 30 AND THE TWISTED IHX LEMMA

This section contains a proof of the “twisted order-raising” Theorem 30 of Section 2, which was used (along with Corollary 5) in Section 3 to prove Theorem 10 of the introduction. A key step in the proof involves a geometric realization of the twisted IHX relation as described in Lemma 42 below.

At the end of the section, the proof of Proposition 14 is completed by Lemma 43 in 4.3 which shows how any order $2n$ twisted Whitney tower can be converted into an order $2n - 1$ framed Whitney tower.

4.1. Proof of Theorem 30. Recall the statement of Theorem 30: If a collection A of properly immersed surfaces in a simply connected 4–manifold supports an order n twisted Whitney tower \mathcal{W} with $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$, then A is regularly homotopic (rel ∂) to A' supporting an order $n + 1$ twisted Whitney tower.

Recall also from subsection 2.8 that the intersection forest $t(\mathcal{W})$ of an order n twisted Whitney tower \mathcal{W} is a disjoint union of signed trees which can be considered to be immersed in \mathcal{W} . The order n trees in $t(\mathcal{W})$ (together with the order $n/2$ ∞ -trees if n is even) represent $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$, and the proof of Theorem 30 involves controlled manipulations of \mathcal{W} which first convert $t(\mathcal{W})$ into “algebraically canceling” pairs of isomorphic trees with opposite signs, and then exchange these for “geometrically canceling” intersection points which are paired by a new layer of Whitney disks. We pause here to clarify these notions:

4.1.1. Algebraic versus geometric cancellation. If Whitney disks W_I and W_J in \mathcal{W} intersect transversely in a pair of points p and p' , then t_p and $t_{p'}$ are isomorphic (as labeled, oriented trees). If p and p' have opposite signs, and if the ambient 4–manifold is simply connected, then there exists a Whitney disk $W_{(I,J)}$ pairing p and p' , and we say that $\{p, p'\}$ is a *geometrically canceling* pair. In this setting we also refer to $\{\epsilon_p \cdot t_p, \epsilon_{p'} \cdot t_{p'}\}$ as a geometrically canceling pair of signed trees in $t(\mathcal{W})$ (regarding them as subsets of \mathcal{W} associated to the geometrically canceling pair of points).

On the other hand, given transverse intersections p and p' in \mathcal{W} with $t_p = t_{p'}$ (as labeled oriented trees) and $\epsilon_p = -\epsilon_{p'}$, we say that $\{p, p'\}$ is an *algebraically canceling* pair of intersections, and similarly call $\{\epsilon_p \cdot t_p, \epsilon_{p'} \cdot t_{p'}\}$ an algebraically canceling pair of signed trees in $t(\mathcal{W})$. Changing the orientations at a *pair* of trivalent vertices in any tree t_p does not change its value in \mathcal{T} by the AS relations, and (as discussed in 3.4 of [25]) such orientation changes can be realized by changing orientations of Whitney disks in \mathcal{W} together with our orientation conventions (2.5).

Any geometrically canceling pair is also an algebraically canceling pair, but the converse is clearly not true as an algebraically canceling pair can have *corresponding trivalent vertices* lying in *different Whitney disks*. A process for converting algebraically canceling pairs into geometrically canceling pairs by manipulations of the Whitney tower is described in 4.5 and 4.8 of [25].

Similarly, if a pair of twisted Whitney disks W_{J_1} and W_{J_2} have isomorphic (unoriented) trees J_1^∞ and J_2^∞ with opposite twistings $\omega(W_{J_1}) = -\omega(W_{J_2})$, then the Whitney disks form an *algebraically canceling* pair (as do the corresponding signed ∞ -trees in $t(\mathcal{W})$). Note that the orientations of the ∞ -trees are not relevant here by the independence of $\omega(W)$ from the orientation of W and the symmetry relations in \mathcal{T}^∞ . A geometric construction for eliminating algebraically canceling pairs of twisted Whitney disks from a twisted Whitney tower will be described below.

Outline of the proof: To motivate the proof of Theorem 30 we summarize here how the methods of [7, 21, 25] (as described in Section 4 of [25]) apply in the framed setting to prove the analogous order-raising theorem in framed setting (Theorem 24 of Section 2): The first part of the proof changes the intersection forest $t(\mathcal{W})$ so that all trees occur in algebraically canceling pairs by using the 4-dimensional IHX construction of [7] to realize IHX relators, and by adjusting Whitney disk orientations as necessary to realize AS relations. The second part of the proof uses the Whitney move IHX construction of [21] to “simplify” the shape of the algebraically canceling pairs of trees. Then the third part of the proof uses controlled homotopies to exchange the simple algebraic canceling pairs for geometrically canceling intersection points which are paired by a new layer of Whitney disks as described in 4.5 of [25].

Extending these methods to the present twisted setting will require two variations: realizing the new relators in \mathcal{T}_n^∞ , and achieving an analogous geometric cancellation for twisted Whitney disks corresponding to algebraically canceling pairs of (simple) ∞ -trees. We will concentrate on these new variations, referring the reader to [7, 21, 25] for the other parts just mentioned.

Notation and conventions:

By Lemma 33 it may be assumed that \mathcal{W} is split at each stage of the constructions throughout the proof, so that all trees in $t(\mathcal{W})$ are embedded in \mathcal{W} .

In spite of modifications, \mathcal{W} will not be renamed during the proof.

Throughout this section we will notate elements of $t(\mathcal{W})$ as formal sums, representing disjoint union by juxtaposition.

Note that if \mathcal{W} is an order n twisted Whitney tower, then the intersection forest $t(\mathcal{W})$ may contain higher order trees and ∞ -trees in addition to those representing $\tau_n^\infty(\mathcal{W})$. These higher-order elements of $t(\mathcal{W})$ can be ignored throughout the proof for the following reasons: On the one hand, in a split \mathcal{W} all the constructions leading to the elimination of unpaired order n intersections (and twisted Whitney disks of order $n/2$)




FIGURE 13. Boundary-twisting a Whitney disk W changes $\omega(W)$ by ± 1 and creates an intersection point with one of the sheets paired by W . The horizontal arcs trace out part of the sheet, the dark non-horizontal arcs trace out the newly twisted part of a collar of W , and the grey arcs indicate part of the Whitney section over W . The bottom-most intersection in the middle picture corresponds to the ± 1 -twisting created by the move.

of \mathcal{W} can be carried out away from any higher-order elements of $t(\mathcal{W})$. Alternatively, one could first exchange all twisted Whitney disks of order greater than $n/2$ for unpaired intersections of order greater than n by boundary-twisting (Figure 13). Then, all intersections of order greater than n can be converted into many algebraically canceling pairs of order n intersections by repeatedly “pushing down” unpaired intersections until they reach the order zero disks, as illustrated for instance in Figure 12 of [22] (assuming, as we may, that \mathcal{W} contains no Whitney disks of order greater than n).

Thus, we can and will assume throughout the proof that $t(\mathcal{W})$ represents $\tau_n^\infty(\mathcal{W})$.

The odd order case: Given \mathcal{W} of order $2n - 1$ with $\tau_{2n-1}^\infty(\mathcal{W}) = 0 \in \mathcal{T}_{2n-1}^\infty$, it will suffice to modify \mathcal{W} — while only creating unpaired intersections of order at least $2n$ and twisted Whitney disks of order at least n — so that all order $2n - 1$ trees in $t(\mathcal{W})$ come in algebraically canceling pairs of trees (since by [25] the corresponding algebraically canceling pairs of order $2n - 1$ intersection points can be exchanged for geometrically canceling intersections which are paired by Whitney disks, as mentioned just above).

Since $\tau_{2n-1}^\infty(\mathcal{W}) = 0 \in \mathcal{T}_{2n-1}^\infty$, the intersection forest $t(\mathcal{W})$ is in the span of IHX and boundary-twist relators, after choosing Whitney disk orientations to realize AS relations as necessary. By locally creating intersection trees of the form $+I - H + X$ using the 4-dimensional geometric IHX theorem of [7] (and by choosing Whitney disk orientations to realize AS relations as needed), \mathcal{W} can be modified so that all order $2n - 1$ trees in $t(\mathcal{W})$ either come in algebraically canceling pairs, or are boundary-relator trees of the form $\pm\langle(i, J), J\rangle$.

For each tree of the form $t_p = \pm\langle(i, J), J\rangle$ we can create an algebraically canceling $t_{p'} = \mp\langle(i, J), J\rangle$ at the cost of only creating order n ∞ -trees as follows. First use Lemma 14 of [25] (Lemma 3.6 of [21]) to move the unpaired intersection point p so that $p \in W_{(i,J)} \cap W_J$. Now, by boundary-twisting $W_{(i,J)}$ into its supporting Whitney disk W'_J (Figure 13), an algebraically canceling intersection $p' \in W_{(i,J)} \cap W'_J$ can be created at the cost of changing the twisting $\omega(W_{(i,J)})$ by ± 1 . Since $\langle(i, J), J\rangle$ has an order 2 symmetry, the canceling sign can always be realized by a Whitney disk orientation choice. This algebraic cancellation of t_p has been achieved at the cost of only adding to $t(\mathcal{W})$ the order n ∞ -tree $(i, J)^\infty$ corresponding to the ± 1 -twisted order n Whitney disk $W_{(i,J)}$.

Having arranged that all the order $2n - 1$ trees in $t(\mathcal{W})$ occur in algebraically canceling pairs, applying the tree-simplification and geometric cancellation described in [25] to all these algebraically canceling pairs yields an order $2n$ twisted Whitney tower \mathcal{W}' .

The even order case: For \mathcal{W} of order $2n$ with $\tau_{2n}^\infty(\mathcal{W}) = 0 \in \mathcal{T}_{2n}^\infty$, we arrange for $t(\mathcal{W})$ to consist of only algebraically canceling pairs of generators by realizing all relators in \mathcal{T}_{2n}^∞ , then construct an order $2n+1$ twisted Whitney tower by introducing a new method for geometrically cancelling the pairs of twisted Whitney disks (while the algebraically canceling pairs of non- ∞ trees lead to geometrically canceling intersections as before):

First of all, the order 0 case corresponding to linking numbers is easily checked, so we will assume $n \geq 1$.

The IHX relators and AS relations for non- ∞ trees can be realized as usual.

Note that any signed tree $\epsilon \cdot J^\infty \in t(\mathcal{W})$ does not depend on the orientation of the tree J because changing the orientation on the corresponding twisted Whitney disk W_J does not change $\omega(W_J)$.

For any rooted tree J the relator $\langle J, J \rangle - 2 \cdot J^\infty$ corresponding to the interior-twist relation can be realized as follows. Use finger moves to create a clean framed Whitney disk W_J . Performing a positive interior twist on W_J as in Figure 14 creates a self-intersection $p \in W_J \cap W_J$ with $t_p = \langle J, J \rangle$ and changes the twisting $\omega(W_J)$ of W_J to -2 . The negative of the relator is similarly constructed starting with a negative twist.

The relator $-I^\infty + H^\infty + X^\infty - \langle H, X \rangle$ corresponding to the twisted IHX relation is realized as follows. For any rooted tree I , create a clean framed Whitney disk W_I by finger moves. Then split this framed Whitney disk using the twisted finger move of Lemma 33 into two clean twisted Whitney disks with twistings $+1$ and -1 , and associated




FIGURE 14. A $+1$ interior twist on a Whitney disk changes the twisting by -2 , as is seen in the pair of -1 intersections between the black vertical slice of the Whitney disk and the grey slice of a Whitney-parallel copy. Note that the pair of (positive) black-grey intersections near the $+1$ intersection is just an artifact of the immersion of the normal bundle into 4-space and does not contribute to the relative Euler number.

signed ∞ -trees $+I^\infty$ and $-I^\infty$, respectively. The next step is to perform a $+1$ -twisted version (described in Lemma 42 below) of the “Whitney move IHX” construction of Lemma 7.2 in [21], which will replace the $+1$ -twisted Whitney disk by two $+1$ -twisted Whitney disks having ∞ -trees $+H^\infty$ and $+X^\infty$, and containing a single negative intersection point with tree $-\langle H, X \rangle$, where H and X differ locally from I as in the usual IHX relation. Thus, any Whitney tower can be modified to create exactly the relator $-I^\infty + H^\infty + X^\infty - \langle H, X \rangle$, for any rooted tree I . The negative of the relator can be similarly realized by using Lemma 42 applied to the -1 -twisted I -shaped Whitney disk.

So since $\tau_{2n}^\infty(\mathcal{W})$ vanishes, it may be arranged, by realizing relators as above, that all the trees in $t(\mathcal{W})$ occur in algebraically canceling pairs. Now, by repeated applications of Lemma 42 below, the algebraically canceling pairs of clean ± 1 -twisted Whitney disks can be exchanged for (many) algebraically canceling pairs of clean ± 1 -twisted Whitney disks, all of whose trees are *simple* (right- or left-normed), with the ∞ -label at one end of the tree as illustrated in Figure 15 – this also creates more algebraically canceling pairs of non- ∞ trees (the “error term” trees in Lemma 42).

As in the odd case, all algebraically canceling pairs of intersections with non- ∞ trees can be exchanged for geometrically canceling pairs by [25]. To finish building the desired order $2n + 1$ twisted Whitney tower, we will describe how to eliminate the remaining algebraically canceling pairs of clean twisted order n Whitney disks (all having simple trees) using a construction that bands together Whitney disks and is additive on twistings. This construction is an iterated elaboration of a construction originally from Chapter 10.8 of [14] (which was used to

show that that $\tau_1 \otimes \mathbb{Z}_2$ did not depend on choices of pairing intersections by Whitney disks).




FIGURE 15. The simple twisted tree J_n^∞ .

Consider an algebraically canceling pair of clean ± 1 -twisted Whitney disks W_{J_n} and W'_{J_n} , whose simple ∞ -trees $+J_n^\infty$ and $-J_n^\infty$ are as in Figure 15, using the notation $J_n = (\cdots((i, j_1), j_2), \cdots, j_n)$. Each trivalent vertex corresponds to a Whitney disk, and we will work from left to right, starting with the order one Whitney disks $W_{(i, j_1)}$ and $W'_{(i, j_1)}$, banding together Whitney disks of the same order from the two trees, while only creating new unpaired intersections of order greater than $2n$. At the last step, W_{J_n} and W'_{J_n} will be banded together into a single framed clean Whitney disk, providing the desired geometric cancellation. (The reason for working with *simple* trees is that the construction for achieving geometric cancellation requires *connected* surfaces for certain steps. For instance, the following construction only gets started because the left most trivalent vertices of an algebraically canceling pair of simple trees correspond to Whitney disks which pair the connected order zero surfaces D_i and D_{j_1} .)

To start the construction consider the Whitney disks $W_{(i, j_1)}$ and $W'_{(i, j_1)}$, pairing intersections between the order zero immersed disks D_i and D_{j_1} . Let V be another Whitney disk for a canceling pair consisting of one point from each of the points paired by $W_{(i, j_1)}$ and $W'_{(i, j_1)}$. Figure 16 illustrates how a parallel copy V' of V can be banded together with $W_{(i, j_1)}$ and $W'_{(i, j_1)}$ to form a Whitney disk $W''_{(i, j_1)}$ for the remaining canceling pair. The twisting of $W''_{(i, j_1)}$ is the sum of the twistings on $W_{(i, j_1)}$, $W'_{(i, j_1)}$, and V ; so $W''_{(i, j_1)}$ is framed if V is framed, since both $W_{(i, j_1)}$ and $W'_{(i, j_1)}$ are framed for $n > 1$ (and in the $n = 1$ case $W_{(i, j_1)} = W_{J_n}$ and $W'_{(i, j_1)} = W'_{J_n}$ contribute canceling ± 1 twistings). If V is both framed and clean, then the result of replacing $W_{(i, j_1)}$ and $W'_{(i, j_1)}$ by V and $W''_{(i, j_1)}$ preserves the order of \mathcal{W} and creates no new intersections.

So if $n = 1$, then W_{J_n} and W'_{J_n} have been geometrically canceled, meaning that their corresponding ∞ -trees have been eliminated from

$t(\mathcal{W})$ without creating any new unpaired order $2n$ intersections or new twisted order n Whitney disks.




FIGURE 16. The Whitney disks $W_{(i,j_1)}$, $W'_{(i,j_1)}$, and V' are banded together to form the Whitney disk $W''_{(i,j_1)}$ pairing the outermost pair of intersections between D_i and D_{j_1} . In the cases $n > 1$, the interior of $W''_{(i,j_1)}$ contains two pairs of canceling intersections with D_{j_2} (which are not shown), and supports the sub-towers consisting of the rest of the higher-order Whitney disks (that were supported by $W_{(i,j_1)}$ and $W'_{(i,j_1)}$) corresponding to the trivalent vertices in both trees $\pm J_n^\infty$.

The next step shows how V can be arranged to be framed and clean, at the cost of only creating intersections of order greater than $2n$: Any twisting $\omega(V)$ can be killed by boundary twisting V into D_{j_1} . Then, using the construction shown in Figure 17, any interior intersection between V and any K -sheet (e.g. an intersection with D_{j_1} from boundary-twisting) can be pushed down into D_i and paired by a thin Whitney disk $W_{(K,i)}$, which in turn has intersections with the D_{j_1} -sheet that can be paired by a Whitney disk $W_{K_1} := W_{((K,i),j_1)}$ made from a Whitney-parallel copy of $W_{(i,j_1)}$. Now, parallel copies of the Whitney disks from the sub-tower supported by $W_{(i,j_1)}$ can be used to build a sub-tower on W_{K_1} : Using the notation $K_{r+1} = (K_r, i)$, for $r = 1, 2, 3, \dots, n$, the Whitney disk $W_{K_{r+1}}$ is built from a Whitney-parallel copy of W_{J_r} , and pairs intersections between W_{K_r} and j_r . Note that the order of each W_{K_r} is at least r . The top order $W_{K_{n+1}}$ inherits the ± 1 -twisting from

W_{J_n} , and has a single interior intersection with tree $\langle K_{n+1}, J_n \rangle$ which is of order at least $2n + 1$.

For multiple intersections between V and various K -sheets this part of the construction can be carried out simultaneously using nested parallel copies of the thin Whitney disk in Figure 17 and more Whitney-parallel copies of the sub-towers described in the previous paragraph.

The result of the construction so far is that the left-most trivalent vertices of the trees $+J_n^\infty$ and $-J_n^\infty$ now correspond to the *same* order 1 Whitney disk $W''_{J_1} = W''_{(i,j_1)}$, at the cost of having created (after splitting-off) a clean twisted Whitney disk of order at least $n + 1$, and an unpaired intersection of order at least $2n + 1$. In particular, this completes the proof for the case $n = 1$.




FIGURE 17.

For the cases $n > 1$, observe that since W''_{J_1} is *connected*, this construction can be repeated, with W''_{J_1} playing the role of D_i , and D_{j_2} playing the role of D_{j_1} , to get a single order 2 Whitney disk W''_{J_2} which corresponds to the second trivalent vertices from the left in both trees $+J_n^\infty$ and $-J_n^\infty$. By iterating the construction, eventually we band together W_{J_n} and W'_{J_n} into a single framed clean Whitney disk at the last step, having created only clean twisted Whitney disks of order at least $n + 1$, and unpaired intersections of order at least $2n + 1$.

4.2. The twisted IHX lemma. The proof of Theorem 30 is completed by the following lemma which describes a twisted geometric IHX relation based on the framed version in Lemma 7.2 of [21].

Lemma 42. *Any split twisted Whitney tower \mathcal{W} containing a clean $+1$ -twisted Whitney disk with signed ∞ -tree $+I^\infty$ can be modified (in a*

neighborhood of the Whitney disks and local order zero sheets corresponding to I^∞) to a twisted Whitney tower \mathcal{W}' such that $t(\mathcal{W}')$ differs from $t(\mathcal{W})$ exactly by replacing $+I^\infty$ with the signed trees $+H^\infty$, $+X^\infty$, and $-\langle H, X \rangle$, where $+I - H + X$ is a Jacobi relator.

Similarly, a clean -1 -twisted Whitney disk with ∞ -tree $-I^\infty$ in $t(\mathcal{W})$ can be replaced by $-H^\infty$, $-X^\infty$, and $+\langle H, X \rangle$ in $t(\mathcal{W}')$.

Proof. Before describing how to adapt the construction and notation of [21] to give a detailed proof of Lemma 42, we explain why the framed relation $+I = +H - X$ leads to the twisted relation $+I^\infty = +H^\infty + X^\infty - \langle H, X \rangle$. In the framed case, a Whitney disk with tree I is replaced by Whitney disks with trees H and X , such that the new Whitney disks are parallel copies of the original using the Whitney framing, and inherit the framing of the original. In order to preserve the trivalent vertex orientations of the trees, the orientation of the H-Whitney disk is the same as the original I-Whitney disk, and the orientation of the X-Whitney disk is the opposite of the I-Whitney disk. Now, if the original I-Whitney disk was $+1$ -twisted, then both the H- and X-Whitney disks will inherit this same $+1$ -twisting, because the twisting – which is a self-intersection number – is independent of the Whitney disk orientation. The H- and X-Whitney disks will also intersect in a single point with sign -1 , since they inherited opposite orientations from the I-Whitney disk. Thus, (after splitting) a twisted Whitney tower can be modified so that a $+I^\infty$ is replaced by exactly $+H^\infty + X^\infty - \langle H, X \rangle$ in the intersection forest. Similarly, a $-I^\infty$ can be replaced exactly by $-H^\infty - X^\infty + \langle H, X \rangle$.

The framed IHX Whitney-move construction is described in detail in [21] (over four pages, including six figures). We describe here how to adapt that construction to the present twisted case, including the relevant modification of notation. Orientation details are not given in [21], but all that needs to be checked is that the X-Whitney disk inherits the opposite orientation as the H-Whitney disk (given that the tree orientations are preserved, and using our negative-corner orientation convention in 2.5 above). In Lemma 7.2 of [21], the “split sub-tower \mathcal{W}_p ” refers to the Whitney disks and order zero sheets containing the tree t_p of an unpaired intersection p in a split Whitney tower \mathcal{W} . In the current setting, a clean $+1$ -twisted Whitney disk W plays the role of p , and the construction will modify \mathcal{W} in a neighborhood of the Whitney disks and order zero sheets containing the ∞ -tree associated to W . In the notation of Figure 18 of [21], the sub-tree of the I-tree denoted by L contains p , so to interpret the entire construction in our case only requires the understanding that this sub-tree contains the ∞ -label sitting in W . (Note that in Figure 18 of [21] the labels I, J, K

and L denote *sub-trees*, and in particular the I -labeled sub-tree should not be confused with the “I-tree” in the IHX relation.)

In the case where the L -labeled sub-tree is order zero, then L is just the ∞ -label, and the upper trivalent vertex of the I-tree in Figure 18 of [21] corresponds to the clean +1-twisted W , with ∞ -tree $((I, J), K)^\infty$. Then the construction, which starts by performing a Whitney move on the framed Whitney disk $W_{(I,J)}$ corresponding to the lower trivalent vertex of the I-tree, yields the +1-twisted H- and X-Whitney disks as discussed in the first paragraph of this proof, with ∞ -trees $(I, (J, K))^\infty$ and $(J, (I, K))^\infty$, and non- ∞ tree $\langle (I, (J, K)), (J, (I, K)) \rangle$ corresponding to the resulting unpaired intersection (created by taking Whitney-parallel copies of the twisted W to form the H- and X-Whitney disks).

In the case where the L -labeled sub-tree is order 1 or greater, then the upper trivalent vertex of the I-tree in Figure 18 of [21] corresponds to a framed Whitney disk, and Whitney-parallel copies of this framed Whitney disk and the other Whitney disks corresponding to the L -labeled sub-tree are also used to construct the sub-towers containing the +1-twisted Whitney disks with H and X ∞ -trees (which will again will lead to a single unpaired intersection as before). \square

4.3. Twisted even and framed odd order Whitney towers. To complete the proof of Proposition 14 in the introduction, the following lemma implies that $\mathbb{W}_{2n}^\infty \subset \mathbb{W}_{2n-1}$:

Lemma 43. *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports an order $2n$ twisted Whitney tower, then A is homotopic (rel ∂) to A' which supports an order $2n-1$ framed Whitney tower.*

Proof. Let \mathcal{W} be any order $2n$ twisted Whitney tower \mathcal{W} supported by A . If \mathcal{W} contains no order n non-trivially twisted Whitney disks, then \mathcal{W} is an order $2n$ framed Whitney tower, hence also is an order $2n-1$ framed Whitney tower. If \mathcal{W} does contain order n non-trivially twisted Whitney disks, they can be eliminated at the cost of only creating intersections of order at least $2n-1$ as follows:

Consider an order n twisted Whitney disk $W_J \subset \mathcal{W}$ with twisting $\omega(W_J) = k \in \mathbb{Z}$. If W_J pairs intersections between an order zero surface A_i and an order $n-1$ Whitney disk W_I then $J = (i, I)$, and by performing $|k|$ boundary-twists of W_J into W_I , W_J can be made to be framed at the cost of only creating $|k|$ order $2n-1$ intersections, whose corresponding trees are of the form $\langle (i, I), I \rangle$.

If W_J pairs intersections between two Whitney disks, then by applying the twisted geometric IHX move of Lemma 42 as necessary, W_J can

be replaced by a union of order n twisted Whitney disks each having a boundary arc on an order zero surface as in the previous case, at the cost of only creating unpaired intersections of order $2n$, each of which is an error term in Lemma 42. \square

5. PROOF OF THEOREM 7

This section defines the doubling map $\Delta_{2n-1} : \mathbb{Z}_2 \otimes \mathcal{T}_{n-1} \rightarrow \mathcal{T}_{2n-1}$ which determines the framing relations described in the introduction, and strengthens the obstruction theory for framed Whitney towers described in [25] by showing that the vanishing of $\tau_{2n-1}(\mathcal{W})$ in the reduced group $\tilde{\mathcal{T}}_{2n-1} := \mathcal{T}_{2n-1}/\text{Im}(\Delta_{2n-1})$ is sufficient for the promotion of \mathcal{W} to a Whitney tower of order $2n$. This means that \mathcal{T}_n can be replaced everywhere by $\tilde{\mathcal{T}}_n$ (with $\tilde{\mathcal{T}}_{2n} := \mathcal{T}_{2n}$) throughout Section 3, and in particular proves Theorem 7 of the introduction.

Definition 44. The *doubling map* $\Delta_{2n-1} : \mathbb{Z}_2 \otimes \mathcal{T}_{n-1} \rightarrow \mathcal{T}_{2n-1}$ is defined for generators $t \in \mathcal{T}_{n-1}$ by

$$\Delta(t) := \sum_{v \in t} \langle i(v), (T_v(t), T_v(t)) \rangle$$

where $T_v(t)$ denotes the rooted tree gotten by replacing v with a root, and the sum is over all univalent vertices of t , with $i(v)$ the original label of the univalent vertex v .

That Δ_{2n-1} is well-defined as a homomorphism on \mathcal{T}_{n-1} is clear since AS and IHX relations go to doubled relations. The image of Δ_{2n-1} is 2-torsion by AS relations and hence it factors through $\mathbb{Z}_2 \otimes \mathcal{T}_{n-1}$. See Figure 18 for explicit illustrations of Δ_1 and Δ_3 .

The following theorem strengthens Theorem 24 in Section 2.

Theorem 45. *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports a framed Whitney tower \mathcal{W} of order $(2n - 1)$ with $\tau_{2n-1}(\mathcal{W}) \in \text{Im}(\Delta_{2n-1})$, then A is homotopic (rel ∂) to A' which supports a framed Whitney tower of order $2n$.*

Proof. As discussed above in the outline the proof of Theorem 30 (just after subsubsection 4.1.1), to prove Theorem 45 it will suffice to show that the intersection forest $t(\mathcal{W})$ can be changed by trees representing any element in $\text{Im}(\Delta_{2n-1}) < \mathcal{T}_{2n-1}$ at the cost of only introducing trees of order greater than or equal to $2n$, so that the order $2n - 1$ trees in $t(\mathcal{W})$ all occur in algebraically canceling pairs. Note that $\text{Im}(\Delta_{2n-1})$ is 2-torsion by the AS relations, so orientations and signs are not an issue here. As in Section 4, elements of $t(\mathcal{W})$ will be denoted by formal sums, and \mathcal{W} will not be renamed as modifications are made.




FIGURE 18. The map $\Delta_{2n-1} : \mathbb{Z}_2 \otimes \mathcal{T}_{n-1} \rightarrow \mathcal{T}_{2n-1}$ in the cases $n = 1$ and $n = 2$.

The case $n = 1$: Given any order zero tree $\langle i, j \rangle$, create a clean framed Whitney disk $W_{(i,j)}$ by performing a finger move between the order zero surfaces A_i and A_j . Then use a twisted finger move (Figure 8) to split $W_{(i,j)}$ into two twisted Whitney disks with associated trees $(i, j)^\infty - (i, j)^\infty$. Now boundary-twist each Whitney disk into a different sheet to recover the framing and add

$$\langle i, (i, j) \rangle + \langle j, (i, j) \rangle = \Delta_1(\langle i, j \rangle)$$

to $t(\mathcal{W})$. Alternatively, after creating the framed $W_{(i,j)}$, perform an interior twist on $W_{(i,j)}$ to get $\omega(W_{(i,j)}) = \pm 2$, then kill $\omega(W_{(i,j)})$ by two boundary-twists, one into each sheet, again adding $\langle i, (i, j) \rangle + \langle j, (i, j) \rangle$ to $t(\mathcal{W})$. Note that $\text{Im } \Delta_1$ in \mathcal{T}_1 corresponds to the order 1 FR framing relation of [23, 24].

The cases $n > 1$: For any order $n-1$ tree $\langle i, (I_1, I_2) \rangle$, create a clean $W_{(i,(I_1,I_2))}$ by finger moves. (Here we are taking any order $n-1$ tree, choosing an i -labeled univalent vertex, and writing it as the inner product of the order zero rooted tree i and the remaining order $n-1$ tree.) Then split $W_{(i,(I_1,I_2))}$ using a twisted finger move to get two twisted Whitney disks each having associated ∞ -tree $(i, (I_1, I_2))^\infty$. Leave one of these twisted Whitney disks alone, and to the other apply the twisted geometric IHX Whitney-move (Lemma 42 of Section 4) to replace $(i, (I_1, I_2))^\infty$ by $(I_1, (I_2, i))^\infty + (I_2, (i, I_1))^\infty - \langle (I_1, (I_2, i)), (I_2, (i, I_1)) \rangle$ in $t(\mathcal{W})$. Note that the tree $\langle (I_1, (I_2, i)), (I_2, (i, I_1)) \rangle$ is order $2n$. If I_1 and I_2 are not both order zero then continue to apply the twisted geometric IHX Whitney-move (pushing the ∞ -labeled vertices away from the ∞ -labeled vertex that is adjacent to the original i -labeled vertex)




FIGURE 19. Multiple ∞ -roots attached to a tree represent sums (disjoint unions) of trees. On the left: the two trees that result from twist-splitting a clean $W_{(i,(I_1,I_2))}$ in the case $\langle i, (I_1, I_2) \rangle = \langle i, ((j, k), (a, (b, c))) \rangle$. Each arrow indicates an application of a twisted IHX Whitney move, which pushes ∞ -roots towards the univalent vertices. The right-most sum of trees becomes the image of $\langle i, (I_1, I_2) \rangle$ under Δ after applying boundary-twists to the associated twisted Whitney disks.

until the resulting union of trees has all ∞ -labeled vertices adjacent to a univalent vertex (all twisted Whitney disks have a boundary arc on an order zero surface) – see Figure 19 for an example. Then, boundary-twisting each twisted Whitney disk into the order zero surface recovers the framing on each Whitney disk and the resulting change in $t(\mathcal{W})$ is a sum of trees as in the right hand side of the equation in Definition 44 representing the image of $\langle i, (I_1, I_2) \rangle$ under Δ_{2n-1} , together with trees of order at least $2n$. \square

REFERENCES

- [1] **J W Cannon**, *The recognition problem: what is a topological manifold?*, Bull. AMS 84 (1978) 832–866.
- [2] **T Cochran**, *Derivatives of links, Milnor’s concordance invariants and Massey products*, Mem. Amer. Math. Soc. Vol. 84 No. 427 (1990).
- [3] **T Cochran**, *k -cobordism for links in S^3* , Trans. Amer. Math. Soc. 327 no 2 (1991) 641–654.
- [4] **T Cochran, K Orr, P Teichner**, *Knot concordance, Whitney towers and L^2 -signatures*, Annals of Math., Volume 157 (2003) 433–519.
- [5] **J Conant, P Teichner**, *Grope cobordism of classical knots*, Topology 43 (2004) 119–156.
- [6] **J Conant, P Teichner**, *Grope Cobordism and Feynman Diagrams*, Math. Annalen 328 (2004) 135–171.
- [7] **J Conant, R Schneiderman, P Teichner**, *Jacobi identities in low-dimensional topology*, Compositio Mathematica 143 Part 3 (2007) 780–810.
- [8] **J Conant, R Schneiderman, P Teichner**, *Higher-order intersections in low-dimensional topology*, Proc. Natl. Acad. Sci. USA 2011 108 (20) 8081–8084

- [9] **J Conant, R Schneiderman, P Teichner**, *Milnor Invariants and Twisted Whitney Towers*, preprint (2010) arXiv:1102.0758v1 [math.GT].
- [10] **J Conant, R Schneiderman, P Teichner**, *Tree homology and a conjecture of Levine*, preprint (2010) <http://arxiv.org/abs/1012.2780>
- [11] **J Conant, R Schneiderman, P Teichner**, *Universal quadratic refinements and untwisting Whitney towers*, preprint (2010) <http://arxiv.org/abs/1101.3480>
- [12] **J Conant, R Schneiderman, P Teichner**, *Geometric filtrations of string links and homology cylinders*, preprint (2012) arXiv:1202.2482v1 [math.GT].
- [13] **J Conant, R Schneiderman, P Teichner**, *Whitney tower concordance of classical links*, preprint (2012) arXiv:1202.3463v1 [math.GT].
- [14] **M Freedman, F Quinn**, *The topology of 4-manifolds*, Princeton Math. Series 39 Princeton, NJ, (1990).
- [15] **M Freedman, P Teichner**, *4-manifold topology II: Dwyer's filtration and surgery kernels*, Invent. Math. 122 (1995) 531–557.
- [16] **M Goussarov**, *On n -equivalence of knots and invariants of finite degree*, Topology of manifolds and varieties, Adv. Soviet Math. 18, Amer. Math. Soc. Providence RI (1994) 173–192.
- [17] **M Goussarov**, *Variations of knotted graphs. The geometric technique of n -equivalence*, (Russian) Algebra i Analiz 12 (2000), no. 4, 79–125; translation in St. Petersburg Math. J. 12 (2001), no. 4, 569–604.
- [18] **K Habiro**, *Claspers and finite type invariants of links*, Geom. Topol. 4 (2000) 1–83.
- [19] **J Levine**, *Addendum and correction to: Homology cylinders: an enlargement of the mapping class group*, Alg. and Geom. Topology 2 (2002) 1197–1204.
- [20] **J Levine**, *Labeled binary planar trees and quasi-Lie algebras*, Alg. and Geom. Topology 6 (2006) 935–948.
- [21] **R Schneiderman**, *Whitney towers and Groves in 4-manifolds*, Trans. Amer. Math. Soc. 358 (2006), 4251–4278.
- [22] **R Schneiderman**, *Simple Whitney towers, half-groves and the Arf invariant of a knot*, Pacific Journal of Mathematics, Vol. 222, No. 1, Nov (2005).
- [23] **R Schneiderman**, *Stable concordance of knots in 3-manifolds*, Alg. and Geom. Topology 10 (2010) 37–432.
- [24] **R Schneiderman, P Teichner**, *Higher order intersection numbers of 2-spheres in 4-manifolds*, Alg. and Geom. Topology 1 (2001) 1–29.
- [25] **R Schneiderman, P Teichner**, *Whitney towers and the Kontsevich integral*, Proceedings of a conference in honor of Andrew Casson, UT Austin 2003, Geometry and Topology Monograph Series, Vol. 7 (2004), 101–134.
- [26] **A. Scorpan**, *The Wild World of 4-Manifolds*, American Mathematical Society (2005).
- [27] **P Teichner**, *Knots, von Neumann Signatures, and Grope Cobordism*. Proceedings of the International Congress of Math. Vol II: Invited Lectures (2002) 437–446.
- [28] **P Teichner**, *What is ... a grope?* Notices of the A.M.S. Vol. 54 no. 8 Sep 2004, 894–895.

DEPT. OF MATHEMATICS, UNIVERSITY OF TENNESSEE, KNOXVILLE, TN

E-mail address: robert.schneiderman@lehman.cuny.edu

DEPT. OF MATHEMATICS AND COMPUTER SCIENCE, LEHMAN COLLEGE, CITY
UNIVERSITY OF NEW YORK, BRONX, NY

E-mail address: teichner@mac.com

DEPT. OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, BERKELEY, CA AND
MAX-PLANCK INSTITUT FÜR MATHEMATIK, BONN, GERMANY